

Communications and Control for Electric Power Systems

The 1991 Report

H. Kirkham

(NASA-CR-192794) COMMUNICATIONS
AND CONTROL FOR ELECTRIC POWER
SYSTEMS Report, 1991 (JPL) 87 p

N93-22477

Unclas

G3/32 0154767

January 15, 1992

Prepared for

Office of Energy Management Systems
United States Department of Energy

Through an agreement with

National Aeronautics and
Space Administration

by

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Communications and Control for Electric Power Systems

The 1991 Report

H. Kirkham

January 15, 1992

Prepared for

Office of Energy Management Systems
United States Department of Energy

Through an agreement with

National Aeronautics and
Space Administration

by

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Prepared by the Jet Propulsion Laboratory, California Institute of Technology, for the U.S. Department of Energy through an agreement with the National Aeronautics and Space Administration.

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof.

This publication reports on work performed under NASA Task RE-152, Amendment 203 and sponsored through DOE/NASA Interagency Agreement No. DE-AI01-79ET 29372 (Mod. A009).

ABSTRACT

IN SECTION 1, a long-term strategy for the integration of new control technologies for power generation and delivery is proposed: the industry would benefit from an evolutionary approach that would adapt to its needs future technologies as well as those that it has so far not heeded. The integrated operation of the entire system, including the distribution system, has been proposed as a future goal. Control of new forms of generation will be important. In transmission, the active control of power flow in response to system operational needs could be incorporated. New control technologies can play an important role. Industry conservatism means that opportunities for research and development in the area of overall system integration and optimal operation can justify government involvement.

In Section 2, the AbNET communication protocols are reviewed, and additions that were made in 1991 are described. In the original network, traffic was controlled by polling at the master station, located at the substation, and routed by a flooding algorithm. In a revised version, the polling and flooding are modified. When a station receives a generalized poll, it responds to the poll over the fiber that brought the poll to the station, and repeats the poll on all the other fibers. While this method is slightly less reliable than a "pure" polling/flooding combination, it can lead to increases in the speed of the network. A software simulation of the AbNET protocols, written to explore the effect of half- versus full-duplex operation, is described. It can demonstrate the benefits of routing by flooding in a damaged network. Hardware developed to allow PCs to demonstrate the protocols in a fiber-based network is described.

In Section 3, the question of interfacing low-energy measurement transducers or instrument transformers is considered. There is presently little or no agreement on what the output of optical CTs should be. Some provide a digital output, others analog. Among the analog ones, some give a low voltage out, others a small current. About the only thing that does seem to be accepted is that low-energy devices will not attempt to reproduce signals like the outputs of the high-energy transformers they replace. The requirements of the interface are reviewed, and a method proposed whereby a standard could be developed.

Appendices deal with the calibration of current transducers (low-energy CTs must be calibrated differently from conventional CTs); with Delta modulation, a simple means of serially encoding the output of an OCT; and with noise shaping, a method of digital signal processing that trades off the number of bits in a digital sample for a higher number of samples.

ACKNOWLEDGMENTS

I GRATEFULLY acknowledge the dedication and effort of Heather Friend and Shannon Jackson of the Jet Propulsion Laboratory. They have contributed significantly to the software simulation and the I/O hardware described in Section 2 of the report.

I also acknowledge the contribution of Bob Phen, my Program Manager at JPL. I thank him for the thoughtful reviews, particularly of Section 1, and the ideas contributed in many interesting discussions throughout the year.

Ole Tønnesen, Professor of High Voltage Engineering, invited me to discuss the question of interfacing optical current transducers at the Technical University of Denmark. He deserves my thanks for giving me the justification to write a reasonably formal report on the topic.

Finally, I would like to thank Dietrich Roesler for his continued support. Dietrich has been our sponsor at DOE for some while now. His guidance and direction of our research continue to be helpful and stimulating.

PREFACE

THIS REPORT, like the one written a year ago, contains material on a number of topics. The report covers long-term prospects for system integration, advances in the AbNET communication protocols, and interfacing low-energy measurement transducers. A word or two about how each of these different parts came to be in this report is in order.

Section 1 is a revised version of a paper written for a planning workshop organized by the U.S. Department of Energy that took place in November. The invitation to contribute required a broad look at the future of the industry, an exercise that I had not found time to indulge in for some while. Perhaps I had fallen into a habit that I think many engineers have, of concentrating on the details and not seeing the big picture. In any event, here was an opportunity to speculate on the long-term outlook for the power industry. The material in the report has the benefit of having been presented and discussed at the workshop, and includes improvements resulting from that process.

Section 2 is a summary of the advances made in the AbNET protocols during 1991. Conceptually, the network is improved in performance by a revised flooding/polling approach. Practically, the hardware for a demonstration in our laboratory is ready. A software simulation of the network was undertaken, in the first place, to study the impact of operating half duplex, instead of the full duplex originally planned. The simulation is very detailed. It became apparent that, with the investment of some effort, it could be used to explore many of the features of the AbNET protocols. This was done, and is described.

I have been considering the question of interfacing low-energy measurement transducers for some while, as chairman of the Fiber Optic Sensors Working Group of the Fiber Optics Subcommittee of the Communications Committee of the Power Engineering Society of IEEE. I pulled together some of my correspondence on the subject for a class on fiber optics that I helped teach at the Technical University of Denmark in the summer. Section 3 of this report builds on that material, and contains much that is new. The appendix on the calibration of optical current transducers is entirely new.

Much of the material is opinion, rather than scientific observation. Rather than write in the conventional passive third person of scientific prose, and thereby give the impression that the weight of science is behind the sentiment expressed, I have elected to write these parts of the report in the first person. This style seems to suit the essay nature of the first and third sections of the report.

HAROLD KIRKHAM

PASADENA, CALIFORNIA
DECEMBER 1991

ABOUT THIS DOCUMENT

THIS DOCUMENT was produced on a personal computer. Design, layout and text input were done using WordPerfect version 5.1. Camera-ready copy was produced by a Hewlett-Packard LaserJet III, using Bitstream scalable fonts. The body type is Dutch Roman, set 12 on 14, and the display type is Swiss. The diagrams were created separately using Freelance Graphics from Lotus.

CONTENTS

ABSTRACT	iii
ACKNOWLEDGMENTS	iv
PREFACE	v
ABOUT THIS DOCUMENT	vi
LIST OF FIGURES	ix
LIST OF TABLES	x
 SECTION 1 LONG-TERM PROSPECTS	
FOR INTEGRATED SYSTEM OPERATION	1
1.1 Overview	3
1.2 Electricity Industry Economics	3
1.3 Review of the Possibilities	8
1.3.1 Distribution and Utilization	8
1.3.2 Transmission	11
1.3.3 Generation	12
1.3.4 Measurement and Control	13
1.3.5 Institutional Impacts	16
1.4 Prospects for the Future	21
1.4.1 Need for Government Involvement	22
1.4.2 Technology Transfer	23
1.5 Summary and Conclusions	25
 SECTION 2 AbNET COMMUNICATIONS NETWORK	
2.1 The AbNET Protocols	27
2.1.1 Open System Interconnection Model	28
2.1.2 Review of Original AbNET Design	29
2.1.3 Additions to AbNET Design	31
2.2 Software Simulator	34
2.2.1 Overview	35
2.2.2 Results	36
2.3 I/O Adapter	37

CONTENTS (Continued)

SECTION 3 INTERFACING MEASUREMENT TRANSDUCERS	39
3.1 Why not retain the 5-A Standard?	40
3.2 The Need for Standards	41
3.3 Timing of Standards	41
3.3.1 The Role of the Market	42
3.4 A Close Look at the Existing Standards	43
3.5 The Black Box Approach	47
3.6 Defining the Black Box	47
3.7 Dynamic Range	48
3.8 Frequency Response	49
3.9 Requirements	50
3.10 Additional Considerations	50
3.11 Digital Interface	52
3.11.1 Sampling	52
3.11.2 Problems with Digital Interface	54
3.11.3 A Modified Black Box	55
3.12 Analog Interface	55
3.12.1 Problems with Analog	56
3.12.2 Frequency Modulation	56
3.13 A Perspective	57
3.14 Figures of Merit	58
3.15 A Process	59
3.16 Concluding Remarks	60
APPENDIX A CALIBRATION OF CURRENT TRANSFORMERS	61
A.1 Applicability to Low-Energy Sensors	64
A.2 Digital Techniques	68
A.3 Concluding Remarks	68
APPENDIX B DELTA MODULATION	69
APPENDIX C NOISE SHAPING	71
REFERENCES	75

LIST OF FIGURES

Figure	Page
1-1 Cost and revenue curves for a price-taking firm	4
1-2 Per unit cost and revenue curves for a price-taking firm	5
1-3 Cost and revenue curves for a natural monopoly	5
2-1 International Standards Organization model of Open System Interconnection	28
2-2 Main features of AbNET in OSI framework	30
2-3 Example of simple network	31
2-4 Representative portion of communication network	34
2-5 2-Master version of communication network	34
2-6 Circuit of fiber I/O adapter	38
3-1 Accuracy requirements for metering service	45
3-2 Accuracy requirements for metering service as a function of current	46
3-3 The Black Box approach to the interface	47
3-4 Possible digital representation of current	53
3-5 Modified Black Box interface	55
3-6 Flowchart showing route to an interface standard	60
A-1 Straightforward method for current transducer calibration	61
A-2 Current transformer accuracy test with current comparator	62
A-3 Current transformer accuracy test with direct-null network	64
A-4 Calibration circuit using current balance	65
A-5 Bridge-resistor method for calibrating optical current transducer	66
A-6 Bridge-CT method for calibrating optical current transducer	67
B.1 Simple delta modulation	69
C.1 Requantization with higher sample rate	72

LIST OF TABLES

Table		Page
2-1	Number of hops required to poll a linear network	32
3-1	Interfaces used in selected EPRI projects	42
3-2	Comparison of Interface Options	51
3-3	Printout of spreadsheet for Figures of Merit	59

SECTION 1

LONG-TERM PROSPECTS FOR INTEGRATED SYSTEM OPERATION

ELECTRIC POWER SYSTEMS have been in existence for over 100 years. During this time, there has been enormous growth in their size and complexity. Yet as we approach the 21st century, the way power systems are operated—both from the technical and the philosophical viewpoint—differs little from the way they were operated in the 19th! Apart from the very occasional use of a dc line, the hardware of a modern power system is only a scaled-up version of the earliest ac power systems. Control of the bulk system is centralized, always involving an operator. The protection system operates autonomously, with no reliance on human intervention. Wherever possible, operation avoids reliance on a communication system.

The industry has been characterized by slow but steady growth of the system. It has been possible to *evolve* new components as they became needed. For example, as the load grew, the economy of scale led to the need to locate generators some distance from the loads. Interconnection at higher voltages was justified, requiring the development of components that could operate at ever higher levels. Ultimately, components with basic impulse insulation levels (BIL) of millions of volts were developed.

A side effect of this approach of developing new solutions out of old has been a reluctance in the power industry to embrace new technology. This conservatism is often justified as prudent, based on reliability and safety concerns. It is a thesis of this report that one result of this kind of growth is that the industry has become stagnant. It may not be an exaggeration to say that compared to some technologies, the power industry is outmoded, that the best available technologies are not being used.

Perhaps sensing this, and aware of the importance of the electricity industry to the nation, the U.S. Department of Energy (DOE) has for some while been funding long-term research in the field of electric power systems. In November 1991 the Department held a workshop aimed at providing guidance for planning the direction of this research. This section of the report is based on a paper prepared for the DOE workshop. It includes most of the material in the original paper, and additional material written since the paper was presented to clarify or expand upon some particular point.

It is clear that there is room for improved application of technology. A software

example was presented at the workshop by Patton, Whittier and Walk (1991), based on a survey (Control Center Survey, 1990) of the largest 300 utilities in the country by EPRI (Electric Power Research Institute). Of the 300 utilities, 30% do not have SCADA (Supervisory Control and Data Acquisition). System operators can have only a limited view of how the power system is performing. Of the 70% that do have SCADA, only a third can perform load flows in the control center. It is unlikely the power systems of the other two-thirds can be operated even close to maximum efficiency. In terms of hardware, one particular example springs to mind. Current measurement systems that have evolved out of century-old iron-cored transformer technology are prone to explosive failure in high-voltage applications. Optical and electronic current transducers that do not have this problem have been available for some years, yet they are hardly being used.

There is a large gulf between the state of the art and the state of the industry, clear evidence of industry conservatism. This is not the optimal response. The operating capability that had evolved as generation and load became geographically separate was unable to prevent instabilities in the system that blacked out a large part of the Northeast—twice within the space of just a few years.

It might seem that these examples are evidence against an evolutionary approach. They are not. In fact the opposite is true. It is argued below that evolutionary thinking is important in any endeavor to look at the long term future of the industry. It is the intent of this section to examine the general principles, and find some illustrative examples. Evolution, as a means of moving forward, can be a very fruitful approach.

The notion of evolution in the biological sense as formulated by Darwin had its origins in the work in economics of Adam Smith. In *The Panda's Thumb* Stephen Jay Gould (1982) wrote that

[Darwin's] theory of natural selection is a creative transfer to biology of Adam Smith's basic argument for a rational economy: the balance and order of nature does not arise from a higher external (divine) control, or from the existence of laws operating directly upon the whole, but from struggle among individuals for their own benefits.

The external control that seemed to Darwin's critics to drive species toward perfection is replaced in biology by something analogous to Smith's "unseen hand."

In the 1990s, evolution seems to be better understood than economics: in a reversal of history, it is therefore suggested that a deliberately evolutionary approach makes sense applied to the marketplace. In either field, economics or biology, the existence of a mechanism for change (technical modification or gene mutation) combined with a mechanism to test the merit of change (market acceptance or survival of the fittest) creates a situation in which improvements tend to survive and prosper. As a result, evolution is a mechanism for adapting to changing circumstances in the multi-variable real world. Planned solutions in this environment are doomed to failure—as shown by the demise of the Soviet Union.

Change, variety and competition are essential ingredients. Evolution proceeds by a sort of shotgun effect: random mutations are introduced into a population. The fittest survive, and represent successful adaptations; the least fit die, and presumably disappear from the gene pool. Thus evolution can create the impression of progress. In the field of power systems, reluctance to change and reluctance to test the merit of changes have com-

bined to suppress innovation. Innovations that have taken place at an increasing pace in many technologies have by and large been ignored in power system control. The relay that caused the 1965 Northeast Blackout differed in no significant respect from the relays in use in, say, 1915. *For this industry, evolution has evidently stopped.* There is the impression of no progress.

If things are to change, if new technologies are to be adopted and adapted, evolution provides a strategy for testing them. It is futile to attempt to predict a 30-year future: the science of *chaos* tells us that a small change in the starting conditions of something as complex as the future of this industry can have far-reaching and unforeseen consequences. Complex non-linear systems are unpredictable in the long run. Consider two examples from the industry's past. How far in advance could the use of digital relays have been foreseen? When did optical measurements seem feasible—and when will they be commonplace?

It is similarly impossible to predict just when some of the technical advances that are bound to affect the industry will be important. The best we can hope to do is to recognize the relevant advances when we see them, and be ready to incorporate them as appropriate. This section of the report is aimed at uncovering some possibilities.

1.1 Overview

Any suggestion for change must recognize that there exists a large investment—both financial and intellectual—in the system we have evolved. This system must be our starting point for further development. The challenge is to chart a course from the present state of the art to a future state improved by greater use of advanced technology.

This section of the report will begin with a review of the reasons (economic) that the industry is so conservative, and then investigate the technology by examining current practice and the R&D possibilities in power system control. The technology will be discussed according to the following breakdown:

- Distribution and utilization
- Transmission
- Generation
- Measurement and control

The section finishes by examining institutional impacts, the prospects for the future, and the appropriateness of government involvement. Technology transfer approaches for the integration of new technology are described.

1.2 Electricity Industry Economics

Textbooks on economics introduce the notion of *market structure* by means of *typical firms* and *perfect competition*. The typical firm long-run unit cost/output curve is U-shaped. At low outputs, startup costs and low utilization of investment mean that there is an economy of scale: higher output would result in lower average costs. At high outputs, there are diseconomies of scale (greater use of input resources may drive up those prices; an extra shift may be needed): long-run average cost increases as output increases. Although it can

be argued that few real firms have cost curves of exactly this shape, the model is a useful tool for explaining the role of the market.


A firm is described as a *price taker* if the prices of its inputs and outputs are determined by forces outside its control. The price taker not only treats costs and prices as fixed, it also regards them as being the result of markets that are in equilibrium. That is, all the other companies that are in competition in the market have already made their managerial decisions, and these will not be affected by decisions made by the price taker.

In some regards, the perfectly competitive market is rather like an infinite bus. This theoretical perfect competition is characterized by a large number of small firms, a homogeneous product, well informed buyers and sellers, and ease of entry and exit from the market. (Actually, easy entry and exit from the market are important only to ensure long-term equilibrium.)

Perfect competition is paradoxical. Since all firms are price takers, and the market is in equilibrium, there is no reason to compete for market share! Inter-firm rivalry has no place under conditions of perfect competition. The reason is simple enough: since the products are homogeneous, innovation in the product brings no benefit. There is nothing to be gained by improving efficiency, either. If one firm were able to increase its profitability by reducing manufacturing costs, the other well-informed firms would follow suit, or other firms would enter the market.

Even if it is unrealistic because it ignores the dynamics of competitive advantage, the idea of perfect competition is useful to examine the question, What should be the output of the firm in order to maximize the profit? Figure 1-1 shows the cost and revenue curves as a function of output for a typical price taker. (The revenue curve is, of course, linear, since the price of the product is fixed. This means, in turn, that the concept of price elasticity has no place here.)

Figure 1-1
Cost and revenue curves
for a price-taking firm




In Figure 1-1, the firm's total cost curve is shown with a reverse-S shape. This is typical of operations that are subject to the law of diminishing returns. Beyond some output, short-run total costs will increase more than in proportion to the output. Obviously, profit is maximized at the output corresponding to the greatest separation between the cost curve and the revenue curve. This is most easily seen by differentiating the data of Figure 1-1 with respect to the output, as shown in Figure 1-2.

Figure 1-2 shows that the point of maximum profit, which can be seen in Figure 1-1 to occur at about 85 units of output, corresponds to the point at which the slope of the total cost curve is equal to the slope of the total revenue curve. The slope of the cost curve is known as the *marginal cost*. Profit is maximized when the marginal cost is equal to the

revenue per unit, or price. (Note also that this curve does not disclose when a firm goes from loss to profit, only when profit is maximized.) The goal of matching marginal cost and marginal revenue becomes a general rule.


Figure 1-2
Per unit cost and revenue curves
for a price-taking firm



In practice, few markets fit the model of perfect competition. Certainly, the electric utility industry does not. It comes closer to the model of a *natural monopoly*. For a natural monopoly, total costs are minimized by having just one producer supplying the entire market. There is always an economy of scale: the U-shaped curve does not apply. Total costs consist of fixed costs (such as capital items) and variable costs that always decrease as output increases. The utility can say, with some justice, that it is inherently wasteful for a competitor to install duplicate lines to serve some fraction of the customers.

Because it can be argued that costs would rise if there were competition, the utility is permitted to operate as a monopoly, but with government regulation. The problem facing the regulators charged with this duty is how to decide the manner in which this regulation should take force. Figure 1-3 shows the standard solution in the U.S.

Figure 1-3
Cost and revenue curves
for a natural monopoly



It will be noticed that these curves include a line for *demand*, which was not the case under perfect competition. As a monopoly, the utility would prefer to operate, as expected, where the marginal cost and the marginal revenue are equal. At this point, the demand (output) is quite small (about 30 units in Figure 1-3) but the price (from the demand curve) is satisfyingly high (a little over 100 units in the Figure). Because the price is so much higher than the cost, and because there seems to be a large unsatisfied demand, this operating point does not suit the power industry regulators. (Under perfectly competitive conditions, this profitability would attract other firms. In a natural monopoly situation, it is argued that since capital investment would have to be duplicated, competition would

simply drive up the price.) A goal of regulation is to come reasonably close to the point at which the demand curve and the marginal cost curve intersect, about 65 units in Figure 1-3. The price and the marginal cost curve are then equal, a much greater demand is being satisfied, and the regulator can feel comfortable.

Unfortunately, the long-run average cost exceeds this price, and the company would eventually sustain a loss. They would prefer a higher price. The intersection of the long-run average cost curve and the demand curve seems like an acceptable compromise, and is usually the target.

Practically speaking, the regulators have no real insight into the shape of the demand curve, nor is the position and shape of the cost curve accurately known. An indirect method of regulation is therefore used. But with so much unknown, how is this achieved?

The answer lies in the *rate of return* the company earns. If the price is above the long-term cost, the company can earn a higher rate of return than might seem reasonable to a zealous regulator, and *vice-versa*. So the goal of regulation is achieved by letting the utility rate of return be similar to the rate of return in other industries.¹ With this solution, the exact shape of the cost and demand curves need not be known.

First, the amount of capital (on which the return is being earned) is calculated. In the electric utility business, this amount is known as the *rate base*. Next, the average rate of return for the economy is estimated. These numbers allow the permitted profit to be calculated. The utility is then instructed to develop a tariff schedule that will create this amount of profit. Fine tuning of the schedule is permitted after the regulations go into effect, since external factors could upset the assumptions underlying the calculations.

The regulation process has defects built into it. If the utility management permits costs to rise, they can be sure profits will not suffer. They have none of the usual incentives of company management. Worse, there is a temptation to expand the rate base to increase total return. Under the usual rules, a new office building contributes to the rate base even though it does not generate electricity. Further, the utility is tempted to behave as a price taker, because the rates are all beyond its control. (And in support of the argument that the industry has been somnolent, remember that price takers have no need of innovation.)

The problem is not simple to solve, and the commissions are faced with an even more difficult problem than the one outlined here because, in the real world, economic factors are not time-invariant. If costs rise faster than rates (and rates can only rise when permitted by the infrequent meetings of the commission) the utility will be short of funds. They might seek to add a "pass-through" clause to the tariff. In an entrepreneurial, competitive market, the right response is often innovation. In the regulated, protected world of the utility, short of money because of regulatory lag, innovation does not stand a chance. It is not difficult to see why this industry has stopped evolving.

Because of these defects, many economists are critical of the utility commissions, and some doubt whether regulation really works to the advantage of the consumer. Ronald Coase, winner of the 1991 Nobel prize for economics, has written extensively on the regulation of industry. He has shifted attention from the operation of regulation under ideal conditions to an examination of regulatory institutions as their structure and effects actually exist. He wrote "It is now generally accepted by all students of the subject that

¹In the recently-privatized electricity industry in Britain, the price—rather than the rate of return—is regulated. It remains to be seen how successful this strategy will be.

most (perhaps almost all) government regulation is anticompetitive and harmful in its effects." (Dolan, 1977)

The industry conservatives have in the past used the Commissions' restrictions and the stockholders' reluctance as excuses for the lack of innovation, but these are not the root cause. It is a simple result of regulated economics. By this observation, it is not intended to attach any implication of wrongdoing to the utilities. It was recently written

They have not been short-changing customers in a deliberate attempt to ramp up profit margins: most are state-owned companies with no incentive to be so aggressive. *They have simply been behaving as monopolies customarily do—shelving (or feigning) avoidable innovation, ducking investment risk wherever possible and keeping a beady eye on their own convenience rather than the customer's.*

(The italics in the quote are mine.) This was in a survey of telecommunications (not electric utilities) in *The Economist* (Survey of Telecommunications, 1991). The discussion was of why many European telephone companies have used fake touch-tone phones (that generate rotary-type pulses instead of the beeps that we hear) to avoid the expense and effort of replacing their central-exchange equipment.

The utility regulation situation is changing. Legislation introduced at the federal level over the last decade or so is beginning to have impact at the state level, where the electric utility regulations are devised and enforced. The first set of legislation, the Public Utilities Regulatory Policies Act (PURPA) of 1978, forced the utilities to accept electric power from generating equipment owned by qualifying facilities. It also obliged the utilities to pay for this energy at a price called the "avoided cost." The intent was to facilitate the connection of small power producers, preferably based on renewable resources. To some extent this has worked, but there have been some unexpected results.

Utility opposition has made it difficult for truly small facilities to connect. Large power producers, typically industries with cogeneration, have the resources to overcome the opposition. Second, PURPA left open the question of what was meant by avoided costs. Interpretations that are applied on a statewide basis, even in states served by several utilities, can be unfair. This means that in New York State, for example, avoided cost is defined to be the same in the Adirondacks as it is in lower Manhattan.

The overall result has a bad side and a good one. The effect has been to put many utilities in a situation where they are obliged to buy power they might not need. And the degree of uncertainty in the utility planning process has become greater. On the other hand, an increasing fraction of the electric energy needs of the country is met by non-utility generation (NUG). According to one forecast, the utilities plan to add 71,600 MW of capacity between 1991 and 2000, and in addition anticipate purchasing 26,000 MW of new NUG power (Annual Forecast, 1991). And PURPA has certainly put downward pressure on the level of avoided costs.

There is now talk of making the transmission system open to all. I noted above the similarity of the model of perfect competition to the infinite bus. Any amount of power can be added to an infinite bus without changing its voltage or frequency. To a small generator, a real transmission system has the same appearance. But the effects are felt on the transmission system. With regulations changed to force utilities to accept power inputs they do not need, they will have to consider power flows that were not anticipated when their bulk systems were built.

Other deregulatory changes might be considered. It may be possible for a customer to contract with a utility in another state, and have his local utility deliver the power, in the way a local telephone company delivers a competitor's long-distance service. Operation in this fashion would represent a radical change in the way the industry operates. The technical and institutional barriers to deregulating to this extent will be explored in greater detail later in this section.

The utilities are facing an era of change and uncertainty. They seem not to like it.

1.3 Review of the Possibilities

If economics and regulation have stifled "market pull" for technological advancement, perhaps it is time to see what is available from the "technology push" point of view. There are many opportunities for technological advance, but this report concentrates on one area, power system control. (The reader might view the outcome as one of the random variations that are the raw material of evolution. Whether the ideas will make sense over, say, the next 30 years has yet to be determined.) The various parts of the power system will be examined from the point of view of real-time control, with emphasis on integrating the operation of all parts of the system. The main elements of real-time control are the ability to

- **monitor** essential parameters,
- **locally process** and
- **communicate** this information to control points,
- **analyze** the data,
- **decide** on a control response,
- **transfer** this decision to remote control devices, and
- **all of this** quickly enough to be useful.

The starting point for discussion is the assumption that performing this real-time control for the entire system can improve both the performance and the economics. In order to accomplish this end, improvement will be needed in one or more of these aspects of control in most parts of the power system. The discussion begins with an excellent example of neglect, from the control point of view: the distribution system.

1.3.1 Distribution and Utilization

Most of the revenue of a utility is obtained through the distribution system, and about half of the capital investment of the company is in this part of the system. 70–75% of all losses and 90% of all outages are attributable to the distribution system. Yet its operation is the least monitored and controlled of all. Further, lack of understanding of the load is ultimately the cause of certain instabilities that have been observed on the transmission system. There are thus both local and global arguments in favor of better knowledge and control of the distribution system.

There are many functions that could be controlled in the distribution system (for a list, see *Distribution Automation*, 1988), but since the distribution system "works" in its present condition, utilities have been unwilling to invest in the necessary monitoring hard-

ware or the communication system needed to implement automation. Anyway, how would they make use of the data? The effect has been that, apart from a few (subsidized) demonstrations of the technology, there are no commercial applications of distribution automation.

The necessary monitoring and communications technology is available today, at reasonable cost. For monitoring and control of the system itself, UHF radio communications, operating with carrier frequencies above 800 MHz, can be used (Holte, 1989). No modifications need be made to the distribution system, other than the addition of the RTU (remote terminal unit) and radio system. Capital costs are probably in the order of \$2000 per RTU site. If two RTUs are installed every kilometer for feeder automation, the cost represents about a 12% increase in the cost of the typical 12-kV distribution circuit. (Applied to the whole power system, this would represent a not insignificant 6% increase in the capital investment of the utility. An advantage of work in the distribution system is that it can usually be performed piecemeal, as need arises or as funds become available.)

If access to more points is needed, for example for meter reading or load management, a fiber optics system can be used. Fiber systems also tend to have very much greater bandwidth than UHF radio systems, and so can support applications that require speed. (The AbNET system discussed in section 2 of this report was specifically designed for the distribution system, and is the ideal candidate.) The cost of the fiber is alone likely to be about \$3000/km, but inexpensive RTUs make the total system cost competitive in high-density applications (Kirkham, Friend, Jackson and Johnston, 1990).

It is still difficult to make a convincing case that the expenditure will be economically justified. While some advocates have claimed that one single function of the set collectively known as distribution automation can pay for the entire system (see, for example, Kendrew and Marks, 1989), such claims seem to be exaggerated (Kirkham, Johnston and Friend, 1989). Both DOE and EPRI have funded work aimed at obtaining realistic cost/benefit data. This approach should be continued.

Since the number of functional applications is large, and there is a continuing need to accommodate changing technology, it is argued that applications can be most effectively developed in the evolutionary, competitive marketplace. But there is still considerable potential here for the government to be involved. The work of developing applications is very much in its infancy. There is a need for standards—in data representations, in interface definitions—that can best be developed with the involvement of the utilities, the professional societies, and the DOE. EPRI, who could be expected to represent the utilities, has stated that standards follow practice, and that EPRI is not in the business of defining standards.² The standards that are needed, by defining what could become product boundaries, would speed development, and help promote competition.

In order for distribution automation to be considered as an integral part of power system operation, its impact on existing centralized planning and control strategies must be considered. First, consider the monitoring aspects of distribution automation. At present, the SCADA system makes available in the control center information about the power outputs of the generators and the power flow on the transmission lines. At a substation that serves a load, the load will be measured, typically as the power flow into a transformer.

²It is not clear how their recent endeavors at promoting the Utility Communication Architecture—UCA—fit this philosophy.

There is sufficient information to perform a load flow calculation. But the load is *modelled* in some way that, although based on experience, may not be accurate. The assumption of constant real and reactive power (as a function of voltage) is usually regarded as conservative. But it may not be. If the voltage on the transmission system is low, perhaps because of equipment failure, the automatic voltage regulators on the load side of the transformers may act to hold the secondary voltage constant. If the high-side voltage is then increased, the load will actually be higher (for the same primary voltage) than before. This kind of interaction between the bulk system controls and the distribution system controls is rather unpredictable, because so little of the action of the low side is monitored. There is an opportunity here for coordinated (global) control.

As an example of an unusual (local) control approach, consider a study in the United Kingdom (Black, Formby and Walker, 1985) that showed that 60% of switching operations on the 11-kV system were for the purpose of isolating other similar equipment for maintenance, and 20% were post-fault operations made to identify the faulted section (rather than to isolate it). Each of these operations required a site visit, and was considered costly and time-consuming. The use of switchgear that could only be operated with the system deenergized, but that could be remotely operated, was estimated to result in a capital savings of about 30% in the replacement of life-expired switchgear. The automation hardware was estimated to cost less than this amount.

Although justified because it is economically attractive, this approach represents a new philosophy of operation: shut down the entire feeder—for a moment—in order to do switching. Once the possibility of a different way of operating the system is accepted, other opportunities will probably become apparent. More complex (non-radial) configurations might be considered, or single-phase switching. Studies of this kind of development may be justifiably government funded, since they are unlikely to be undertaken otherwise.

Utilization has been scarcely mentioned in this report so far. Since the concentration is on control, the discussion will be restricted to consideration of ways we can control utilization to be of value in operating the power system, and ways we can affect utilization to benefit the customer.

In the category of improving power system operation, of course, are the various kinds of demand-side management (DSM). Some customers are willing to be denied the use of some of their loads for some of the time, in return for lower rates. The advantage to the utility is the ability to "shave the peak" load and thereby defer capital expenditures. (Note that peak shaving can apply both to system peaks and local load peaks.) As useful as this is, it does not go far enough, and it is usually limited by the communications system to very coarse control.

Typically, controlled loads are turned off for so many minutes in the hour. A better approach, and one that avoids problems when many loads are turned on together, would be to change the set-point of controlled loads. The temperature settings on air-conditioner thermostats could be increased when it is necessary to reduce load. This requires a more sophisticated control and communication system than is presently used for DSM, but it fits with a future in which loads are increasingly automated. (One might note in passing that an opportunity that seems largely to have been missed is to use brute-force load management as an adjunct to load-shedding for stability control.)

The other side of DSM is the improvement of the service offered to the customer. With improved monitoring, it would be possible to inform the customer of not only the amount of electricity used during the billing period, but also the appliances that used it.

This sort of information has not been made available before, so its uses are speculative. But would it not be useful to know, for example, whether the refrigerator were suddenly using more energy? Is it really economical to buy that half-a-sheep, and pay the energy bills for the deep-freeze, or would it be cheaper to buy the meat from the butcher? One of the conservation measures frequently advocated is the use of compact fluorescent lights, instead of incandescent. It should be possible to show the savings to the customer, in a convincing way. Likely, the word would spread, making this kind of conservation more widely practiced.

For decades the utility has provided electric energy with its frequency sufficiently well-controlled that it can be used for time-keeping. This accuracy arises out of the fact that frequency is used as an indication of the balance between generation and load. No extra charge has ever been made for this feature, and most customers have come to count on it.

With the increasing use of solid-state clocks, it becomes possible for a utility to offer additional services. For a small charge, clocks can be reset following power failures, for example. Using the power-line secondary as a local area network, the utility could offer security services: lights could be turned on and off, and curtains moved to simulate occupancy of a house; load monitoring could detect after-hours use of commercial or industrial facilities.

The utility of the future might think of itself not as a supplier of energy, but as a supplier of services. With adequate communications, and increasing use of automatic control, there may be a large number of revenue-earning services that can be offered.

1.3.2 Transmission

Per unit of power handled, the transmission system is the least capital-intensive part of the power system. Over the years, the cost of transmission line losses has been traded-off against the capital cost of improving efficiency, and effective computer programs for optimal line design are available.

The basic problem facing both operator and planner these days is to get more power down a given corridor, without sacrificing reliability or efficiency. For cost reasons, most transmission lines are ac, and for historical reasons they are mostly three-phase. Techniques for increasing the capacity of a line include adding compensation, reconductoring, converting to a higher phase-order and converting to dc. All these approaches are well understood (some of them having been the beneficiaries of past government funding), and research in them is scarcely needed.

New methods of operating the transmission system may be needed if the "open access" policy that is now being discussed becomes a reality (see Section 1.3.5 below). Existing transmission systems have generally been designed to solve very specific problems, delivering power from the utility's generators to the utility's loads. The use of this kind of system as a general purpose power-wheeling facility is bound to cause difficulties both in planning reinforcements and extensions, and in operation.

More widespread use of dc might help. The concept of adding a parallel dc circuit with active stability control is over two decades old (Dougherty and Kirkham, 1970). That work was based on the notion of adding a derivative term to the equations of motion of the power system. (Normally, a power system is characterized by a set of non-linear second

order equations in which the coefficients of the first order terms are zero.) However, the idea of adding a new dc line sized and controlled specifically to stabilize the parallel ac system (rather than to carry power itself), may be new.³

For reasons of transient stability, it is preferred that power lines are operated with power-angles less than 30° . This means that the power transmitted is limited to $0.5 (\sin 30^\circ)$ of the maximum power that *could* be transmitted. In fact, many lines are loaded to a greater extent than this, though not without a certain nervousness. If the operating angle is increased to (say) 45° the transmitted power is increased by 40%, but the system is closer to the stability limit. It would be desirable to design a control to allow this increase in power transfer without its concomitant loss of stability margin. In essence, there would be an increase in the transient stability limit, at the cost of some control. The concept is at least worthy of study. The first question is how large must the dc system be in order not to degrade the stability of the system, and how much would this cost compared to reinforcing the ac system. The second, and more difficult question is how to control such a complex non-linear system.

1.3.3 Generation

In most utilities, power is generated by a mix of energy sources. The control (scheduling, dispatching) of these is well understood. An increasing amount of non-utility generation (NUG) is now being connected to the utilities' power delivery system. Section 210 of the Public Utilities Regulatory Policy Act (PURPA) encourages renewable resource generation, and obliges utilities to buy the power produced by Qualifying Facilities (QFs). There is also a competitive bidding process whereby utilities can invite Independent Power Producers (IPPs) to build generation facilities in a competitive fashion.

The DOE has already invested considerably in R&D on the integration of new-technology generation. Much of that work, which was done in the last decade, was concerned with small generators and relatively low penetration into the power system. Recently there has been a trend for quite large cogeneration facilities to connect to utilities' networks, in some instances requiring the construction of new transmission lines. It may be worthwhile to review the earlier work with a view to its applicability to large IPPs.

Since non-utility generation is a relatively recent phenomenon, brought about by regulatory action, the industry has not had time yet to evolve a strategy for dealing with it. Some utilities, obligated by the terms of PURPA 210, feel that the law benefits only the QF operator. Apart from the question of proper pricing for "avoided costs," it seems reasonable that large IPPs should be integrated into the utility planning and operation strategies to a greater extent than is now required by the law. Integrated control of the system may make it necessary in the future for the utility to be able to exercise more control (and more flexible control) over QFs than is presently usual.

³There are analogies in aerospace. At least one Soviet fighter has been observed at the Paris Air Show to be capable of pulling up its nose and increasing the angle of attack on the wings to the point of stall and beyond (so far, in fact, that the tail of the aircraft is actually ahead of the nose) while maintaining level flight. And the well-known Harrier vertical take-off plane is known to be unstable without its computer control system.

1.3.4 Measurement and Control

Real-time control involves the communication of information to one or more control points, where control action is decided upon. The first part of this is the measurement itself.

It was noted earlier that current transformers built for high voltage use had problems. In defence of the evolutionary approach by which such devices were developed, it is pointed out that evolution offers no guarantee that useful adaptations will continue to be adaptive in changed circumstances. What works well at 500 V may not work at 500 kV.

The world of power system measurements could be revolutionized by the advent of optical measurements. Several companies (in the U.S., Europe and Japan) have independently developed optical transducers to replace current transformers. For historical reasons, these have become known as OCTs, which could stand for optical current transducers or optical current transformers. Some use bulk optics, other fiber optics. All have the advantage that the insulation system of the measurement is greatly simplified. The optical materials used are inherently insulators, and moderately inexpensive. Economically and technically they are better than the conventional CT. Yet in the U.S., apart from a few demonstrations, OCTs sit unused on the factory shelf. While the lack of an interface standard is cited in explanation, see Section 3, the fact is that industry conservatism is at work.

Optics can also be used for distributed measurements. By means of a device known as an Optical Time-Domain Reflectometer (OTDR) the temperature profile of a transmission line can be obtained. This information can be of use in maximizing the power on a transmission line, especially in hot weather when the cooling effect of the air is small, and when many utilities are at peak load. OTDR temperature measurement systems are commercially available. Perhaps research into automating the use of the data would help their more widespread application.

As far as measurements are concerned, it is fair to say that all the parameters needed to operate a power system can be measured with sufficient accuracy and speed. Now we turn to the purpose for which most of this information is obtained: system control as implemented in the energy management center.

The control system can be broken down into several parts: the measurement itself, any local processing, communications (typically SCADA), and energy management center processing. Because of the complexity and size of the typical power system, a large number of data are brought into the energy management center. So much information is delivered that it is easy to overload the system operator. Indeed, this is one of the well-known problems. (It could be worse. Some of the developments suggested below will increase the amount of data considerably.) But operator overload is already a problem well on its way to solution. As more functions are automated, and alarm-handling improved, the overload experienced by operators will inevitably diminish. But system performance will not have improved, without some additional steps being taken.

The software in the energy management center has developed in an application-oriented environment. Often, the software for each function has been developed on its own, with its own user interface, means of accessing and displaying data, and its own model of the system. In the future, the trend will be more information oriented, with the access, display and models increasingly decoupled from the application.

The amount of data acquired every scan of a representative SCADA system, typically every 3 s, is somewhat less than 100 kb, depending on the size of the system and the

parameters monitored. This is about as fast as the computer software can handle it. If the details of the distribution system are added to this data stream, a total in the order of 500 kb/s will be delivered to the energy management center. Many of the data will not change significantly between scans. (In the past this has led to the use of compression techniques.) Some of the data will contain errors.

This information is used to supply a static state-estimator program which rejects bad measurements and, in turn, provides a data-base for the programs controlling power generation, executing voltage control, and performing security assessment. The availability of this common database lends itself to an information-oriented future. It is reasonable to ask what would be done if the limit imposed by communications or computation were removed.

A power system is operated against a constraint of some sort (economic operation, environmental dispatch or capacity limit) all the time. Modern control theory indicates that this is an optimization problem. Optimal power flow programs are presently available. The improvement possible here from the more *frequent* acquisition of data is likely to be small because the state of the system can change only slightly between scans.

More *widespread* data acquisition may have some advantages. A new measurement being considered in some utilities is that of phase angle across the transmission system. (Since this uses the Global Positioning Satellite system for its ultra-precise timing, it is evidence that not all utilities are conservative all the time.)

System phase represents only a small amount of information, used to improve control of the generation and transmission system. The acquisition of additional data, particularly from the distribution system, should be considered. It was observed above that ultimately this was the only way the behavior of the bulk system could be understood. But there are local advantages, too. While power system losses are not large (a few percent), they do mostly (70–75%) occur in the distribution system. Lee and Brooks (1988) used data from Pennsylvania Power and Light obtained throughout 1984 and 1985 at 15-minute intervals to study the effect of system reconfiguration. Losses were reduced over 14% (2500 MWh) with an annual estimated savings of over \$110,000 for 26 feeders. The overall efficiency of the power system was calculated to rise by 0.2% to 98.95%. If these figures could be extrapolated nation-wide, this seemingly small amount could be significant.

It is hard to say what the result of *better* measurements (in terms of bandwidth or accuracy) would be. Improved measurements have made a difference in the past. Here are two examples. First, most of us were taught that when a short circuit is applied at the terminals of a generator, it speeds up. In fact, careful measurements have shown that, momentarily, the generator slows down. Fortunately, the error is on the safe side for relay settings. Second, careful measurements of lightning have shown that the simple steep-rise, slow decay waveform that is used for lightning impulse testing is wrong. Lightning is not only very much faster than our test waveforms, it can be bipolar. In this case, the error is *not* on the safe side for insulation coordination.

In the area of system computation, improvement is also possible. It may be unnecessary for the entire system to be monitored either more rapidly or more widely for more advanced techniques such as dynamic state estimation to be used for transient stability control. (Dynamic state estimation involves the solution of the dynamic equations of the system rather than just the static equations.) With a 3-second scan time for the SCADA, and (say) a 3-ms scan time for certain crucial data, system operation decisions concerning state trajectories could be made rapidly enough to operate 1-cycle breakers, for example,

or for fast-valving of generators. This would require that the data be processed in a few ms, a task that cannot be done with today's techniques. It was observed above that just the static solution took about 3 s.

As another example, consider security assessment. As presently performed, this is a straightforward problem made large by the number of *combinations* of contingencies possible from a given starting point. In fact, the problem becomes so large that conventional techniques cannot solve the problem in real-time.

The technology of artificial neural networks has been proposed as a way of overcoming the problems of solving complex systems rapidly, and some progress has been made in this direction in the field of power systems. Rule-based artificial intelligence (expert systems) has been considered for power system use for some while. The application of neural nets is more recent. A paper by Pao (1990) surveys the application of neural networks to power systems. He shows that neural nets are less memory-intensive than look-up tables, and faster than ordinary computation or other more conventional artificial intelligence techniques. There are hardware problems associated with neural nets, however, and these are discussed below.

Continuing with the discussion of software issues, there are a number of other new developments outside the field of power system analysis that should be considered. *Fuzzy logic control* offers the potential for robust problem solving. The name derives from fuzzy set theory, in which *fuzzy* is used to describe problems containing a degree of imprecision about what is known. Fuzzy logic control seems to be capable of successful operation over a wide range of conditions in the power system, whereas the performance of conventional control would deteriorate because of the system non-linearity.

The relatively new field of *genetic algorithms* has met with success in some areas, including power systems (Ajjarapu and Albanna, 1991; Nara, Shiose, Kitagawa and Ishihara, 1991). Genetic algorithms are optimization and learning techniques based on the ideas of natural selection and genetics. The approach may become useful as a planning or design tool. Control applications may become attractive as neural or parallel processing becomes easier to implement. Applied to neural nets, genetic algorithms offer the prospect of faster learning than conventional methods, thereby making some neural net methods feasible that would otherwise suffer from the combinatorial explosion problem in the training phase.⁴

The discussion of the software issue is closed with a few words of a general nature. There has been a regrettable tendency to avoid "open systems" in power system software. Suppliers have tended to keep their products proprietary. Some utilities that have the expertise to develop their own energy management systems have done so. American Electric Power (AEP), the largest investor-owned utility, has traditionally built its own energy management systems (EMS), for example, and Commonwealth Edison (CECo) has recently followed suit. AEP wanted to have firm control over the EMS design. CECo was attempting to upgrade a system that had, over its lifetime, been more improved by CECo than its original vendor (Hansen, Illian and Kilar, 1991).

The problem with all of this is that closed systems are developed. The possibility of integrating software from other sources is lost. Any new piece of software would have to be adapted for the particular case. Software maintenance becomes more difficult (which

⁴In view of the comments in the Introduction, it is to be expected that I would support study of genetic algorithms.

is the reason why CECo began writing their own software in the first place!), and more costly. It would be to the benefit of all if this trend were reversed.

Open systems are systems that are adaptable, upgradable, can make use of industry standards, and have carefully defined interfaces. Such systems permit a modular approach to software (and hardware, for that matter), and encourage competition in product development. Closed systems are the opposite—frozen in design and function, and headed for obsolescence as soon as they are built.

There are moves towards open systems today. Any such moves, including standardization of interfaces (such as for application/environment, human/machine or communications), are to be encouraged.

Turning now to hardware, the dramatic improvements in performance of computer hardware cannot have escaped anyone's notice. The power industry does not represent a large enough market to drive computer hardware technology, so we are limited to consideration of the advances that have occurred because of other driving forces: optical storage, parallel processing, networking and so on.

There is one more specific observation to be made in this area. It was noted above that one of the problems with the neural net approach is the difficulty of implementing the artificial neural network in hardware—much of the work to date has been done using simulations of networks. Some special-purpose integrated circuit neural networks have been built—JPL, for example, has built some analog networks using very large scale integration—but they have limited resolution and are rather expensive.

A new method of handling analog data has emerged recently: digital signal processing. Integrated circuits designed specifically for DSP use pipelining techniques and parallel processing. They have shown themselves to be capable of impressive performance in a number of applications. For example, when the "pause" button is pressed on a CD player, the disc keeps rotating, but the sound stops. The signal is not interrupted, rather the 16-bit signal samples are multiplied by a number series that represents a rapidly decreasing function. In order to avoid an audible click, a DSP chip is performing calculations at a rate in the Mflop region!

To achieve this kind of performance, the internal architecture of the DSP circuit is highly parallel. Whether applied as artificial neural networks, or in some other way, DSP techniques have great potential to solve the problems caused by the complexity or speed requirements of power system control. DSP techniques could have significant impact on protection, too. While early (analog) attempts to use electronics for protection met with difficulty, digital relaying is now becoming accepted. Using DSP, a 1024-point FFT (fast Fourier transform) can be performed in less than 1 ms—orders of magnitude faster than the mechanical time constants in a power system. Using DSPs, more complex calculations than are presently possible in the power system could be done with time to spare, both in relays and in the control center. High-speed correlations, for example, might offer new ways to perform fault identification and location in real time. The possibilities are quite fascinating.

1.3.5 Institutional Impacts

The National Energy Strategy (1991) points out that 45 GW of new capacity could be saved over the next two decades by institutionalizing integrated resource planning (IRP). An inte-

grated approach to operation can result in improved energy use. However, any attempt to integrate the planning or operational control of a power system will be met with one particular problem. Very few utilities are organized in a way that permits planning or operation as a single entity. Almost without exception, generation, transmission and distribution are operated by separate parts of the company, sometimes by different companies. Communications is often handled separately; so is protection.

This kind of organization has resulted in difficulties already. For some while, it has been technically possible to include a fiber optic communication channel in the static wire of a transmission circuit. The method is known as OPGW, optical ground wire. The approach was first used by the Central Electricity Generating Board (CEGB) in England, in about 1980. Yet OPGW was incorporated in Mexico before it was in the U.S., largely because of the difficulties of inducing the transmission department and the communications department of any U.S. utility to work together.

This is not intended to be a finger-pointing exercise. The utilities are not the only source of institutional problems. NUGs are, by definition, separate companies producing electricity. The utilities are obliged to buy QF energy, though they have no say in the locating of the generator, or even whether it is needed. Subtransmission or even transmission system reinforcements might be needed to accommodate a large IPP. If the power output is large enough, planning and operation of the generation facility clearly have to be coordinated with the remainder of the utility. This should include consideration of the utility's estimates for future load growth, which today it typically does not.

The problems of integrating QFs can be expected to become more acute as the proposed reform of the Public Utility Company Holding Act is enacted, and as PURPA is changed to increase the size and number of QFs, all within the aims of the National Energy Strategy. The interconnected transmission system is to become the wholesale marketplace for buying and selling power, no matter the source. While this will have the effect of increasing competition, it will cause some problems.

It is not clear that technical research, government funded or not, can address any of these institutional problems. They are brought up here to indicate that even the perfect technical solution to a control problem cannot be guaranteed successful implementation. Perhaps one way to help is to consider a situation that, by taking an extreme position, makes the likely future seem more reasonable. Let us indulge in that exercise.

Possibly the most dramatic change facing the industry at present is the "open access" to transmission that is being mandated by regulators to encourage wholesale competition. I would like to consider an expanded version of this, hinted at earlier: operating the electric utilities in a considerably deregulated manner, similar to the telephone companies.⁵

Suppose a customer were able to buy electricity from any utility, anywhere. This would surely put the utilities in competition with one-another. The local utility would be responsible only for *delivery*. Guaranteed revenue only from this service, they could make a living out of wheeling charges. Energy sales could be added if they had generation that

⁵When I started to write the description that follows, I thought it had evolved from a concept discussed by T. Paul Mauldin of PG&E in an after-hours conversation at the DOE workshop. His idea was to consider a large substation to be a "profit center," capable of negotiating its own deals for power with the generating stations it was interconnected with (presumed to be owned by the same company). I see now that it also owes much to the deregulation of the phone companies.

was competitive. The cost of the energy delivered to the customer would thus automatically include the cost of generation, paid to the generating company, and the cost of transportation, paid to the local company. As in the case of the telephone companies, the local company would handle all the billing.

One cannot but think that a market-like system of this kind, populated by informed buyers and sellers, would result in a close-to-optimum situation. So what might be the problems or objections?

- First, there is the cost of the transactions needed to make the approach work. Whatever the size of the customer, neither the utility nor the customer can afford to waste resources negotiating.

Just what resources could be used are very technology-dependent. With little change in hardware on the part of the customer, domestic users might be prepared to negotiate a price for the next quarter or year of service. A large industrial customer might wish to review his options more frequently, say every month. With some advances in the hardware installed at the customers' premises, the negotiations could be handled automatically, on a minute-to-minute basis. This would require some kind of computer to represent the customer, and a communication system that—for the present at least—does not exist. Nevertheless, the fact is that the telephone companies have demonstrated that this kind of operation is feasible. The local phone company, a natural monopoly no less than the power company, delivers the long-distance service of its competitors, and bills the customer on their behalf. For many customers, the phone bill is smaller than the electricity bill; yet the billing calculations can be done on a second-to-second basis. It seems that the transaction cost can be made reasonable.

- Second, the utility might object that there would be unexpected or unpredicted power flows. The concern arises because, with ac transmission systems, the power flows according to the basic equations of electrical networks, and is no respecter of ownership.

If unanticipated flows do occur, it is evidence only that the present situation is not optimal. Consider the case of the west coast⁶. Power is generated near population centers all along the coast, but there is an excess of power (from hydro resources) in the north, and a net import into the south. The Pacific Intertie, a combination of ac and dc lines, is used to deliver a large fraction of that power. An unexpected power flow might occur if a large number of customers in, say, the Los Angeles area decided to buy their power from Mexico, with whom the area is connected through San Diego Gas and Electric (SDG&E). Suppose this happened. The unavoidable conclusion is that the power is being bought from the south because the price of Mexican energy was attractive. If the resulting power flow is unexpected, would it not be fair to say that the Los Angeles Department of Water and Power (LADWP) or SDG&E should have been buying their (wholesale) power from Mexico? It is of course most likely that there would be no unexpected power flows. The power systems are not operated a long way from optimum.

- A third difficulty arises when system extensions are considered. The example of the Pacific Intertie can be used again. Suppose the clock is wound back to a time before the Pacific Intertie was built. LADWP is faced with good load growth but environmental

⁶The examples to follow are for the purposes of illustration only. By naming lines and power companies, I have no desire to commend or to slight. It simply happens that I live in the area served by these entities.

problems with building local generation. The Pacific northwest has an excess of generation, but there are no adequate tie-lines to the south.⁷ The question then is, How is the Intertie to be financed?

In reality, the question has already been answered in one way. Bonneville Power Administration operates one of the nation's largest transmission networks. They own no generation, and they sell wholesale. In the case of BPA, the financing was through the use of federal funds, but in principle the need for transmission could have been met by an entrepreneur. Of course, that entrepreneur could have been a business unit of LADWP or the Corps of Engineers (owners of much of the hydro generation). So the question of financing system extensions or reinforcements is really moot. In a competitive, market-driven environment, needs of this kind will be met simply because it is in somebody's financial interest to do so. That this can happen without government subsidies is shown by an investor-owned Swiss company, *Elektrizitätsgesellschaft Laufenburg*, that operates a transmission system connecting Switzerland, Germany and France. Their revenue is earned from carrying charges associated with the power flows needed to take advantage of the diversity of the load in those countries.

- Finally, there are questions of system reliability, stability and security. These are aspects of system operation that require centralized decision-making. Will it be possible to achieve the appropriate control in the decentralized, transaction-driven system we have been describing?

The answer again must be Yes, but there are likely to be some awkward moments. Even in today's protected environment, some utilities are reluctant to furnish system data to their neighbors. These data are needed for accurate calculation of system performance, but are withheld in a spirit of "beggar-my-neighbor." This attitude would not improve if neighbors become competitors.

Perhaps it will be necessary to understand how to cost investments that are made for stability or security, so that these can be factored into the billing process. And if the power flows do vary more widely than before, it may be necessary to implement novel protection schemes, perhaps with non-constant parameters. But surely these are the kind of issues that are in a "grey area" now, and would be clarified only if it became necessary to provide a proper return on the particular investment, rather than the general rate base.

A utility can ascertain what value a customer places on reliability by asking directly, How much would you pay to have the number (or duration) of outages that you experience reduced by some (specified) amount? Have they asked themselves the same question?⁸

⁷I have no idea if this was the situation on the west coast when the Intertie was being considered, or whether the generation followed the transmission, or whether the two went hand in hand. For my present argument, the most realistic case is made if this particular version of the problem can be adequately answered.

⁸They had not in 1970. An answer to this question would bias the "break-even distance" approach to comparing ac and dc, then being used to justify the use of dc transmission. Because of the higher terminal costs and lower line costs, there is a distance above which dc transmission costs less to build. There was—and, as far as I know, is—no consensus on what value to attach to the stability improvement in the ac system achieved by modulating the dc line.

Having argued that there are no insuperable problems, let us turn to examine the advantages of this kind of operation. It was argued above that it was unlikely that there would be any major changes in the power flows in the system. So why bother?

- First, while there might be no abrupt changes in system operation, gradual changes that reflect customer preferences may be expected. This means that the utilities will begin to feel the effect of competition, and respond accordingly. Further, an increased degree of specialization is possible, with generation and transmission being separate functions, so increased efficiency should follow.

In other words, reinforcing the theme of this part of the report, here is a mechanism to put evolution back to work in this industry. It is this change that, in the long run, can be expected to have the most profound beneficial effect.

- Second, the more specific costing should result in a more equitable handling of a number of cost-related factors. In a regulated environment, it is difficult to see how to handle conservation, and because of this relatively little has been done in the area. It is a fundamental precept of market economies that *price* is a conveyor of information. What better way, then, to see what value the customer places on his need for air-conditioning, for example, than to institute load management? System reliability could be handled in a more informed way, too.

Improved cost allocation methods should help deal with environmental matters, too. Under the 1990 Clean Air Act, power generation will have to reduce its emission of sulphur dioxide by more than 50% by 2000. The Act leaves it up to the utilities to decide how the cuts should be made, recognizing the differences that exist between power plants. A system of "pollution permits" has been created, that should allow a market-like process to optimize the pollution reduction. The true cost of reducing air pollution, for example, by retiring a particular generator, rather than adding scrubbers, should become apparent⁹. The value placed on visual amenity should become evident too. The transmission system could be reinforced without significant cost, in an effort to make more use of distant generation. But whether this would be acceptable is a different matter: the transmission system is the most visible part of the power system.

The reader might be tempted at this point to ask, If competition is such a good idea, why has it not happened? Apart from the conservative nature of the industry and the regulators, about which so much has already been said, it would be fair to say that the technology to implement the kind of system we have been considering has not been available much before now. There is an implicit dependence on a communications system, and the widespread availability of computers.

The calculations that must be made for fair and accurate billing of delivery services are perhaps more complex than those needed by the local phone company. In a theoretical,

⁹The interference of local regulators may prevent the market's development. Since the Commissions are state bodies, they can—and do—disregard federal intentions to promote their own interests. Commissions in coal-producing states (such as Kentucky, Ohio and West Virginia) are threatening to adopt policies that will encourage the continued use of coal. This will certainly distort the market. An argument can be made, it seems, for federal regulation to take precedence over state regulation.

perfect system, the capital costs associated with transmission facilities as well as the losses associated with the instantaneous power flows would be accounted for by measurement and calculation, and included in the bill sent to the customer. The bill would thus be dependent on the results of a load flow calculation, which seems unrealistic. In a less precise but probably more practical system, some of these calculations are based on statistics, essentially describing the marginal costs of meeting various loads. A simplified approach, perhaps based on the B-coefficients (loss formula coefficients) described by Kron (1951) and Kirchmayer (1958) might allow for distributed accounting for transmission losses. Refinements could follow if needed.

The fact is, of course, that the billing methods used at present are statistical, incorporating capital costs and energy costs in some broad-brush way that is quite transparent to the customer, but are in any case approximate. A free-market approach simply requires that the various components of cost be known fairly well.

With a competitive system like this, can the regulator be entirely removed from the process? If a convincing case can be made for a Yes answer, I suggest that serious thought be given to investigating the obstacles to implementing deregulation. It is more likely that the answer must be No. The main reason that regulation might have to be retained is that the local company is still a natural monopoly. Because the distribution system is the most costly in terms of investment per customer, or per kW of load, it still seems unrealistic to expect competition at the customer level to attract new investors. A new delivery-only utility would have to build new distribution facilities, and could expect to pick up at most half the existing load. There is also the possibility of the generation and transmission companies forming a cartel, though anti-trust regulations already exist to make this option unattractive. There are still some benefits to be gained from further consideration of a reduced degree of regulation and increased competition, and this possibility should be investigated. At the very least, a study of the impact of this kind of competition might help the regulators see how far we are from optimal operation at present.

1.4 Prospects for the Future

For some time we have been reading predictions of future problems in the utility industry. There would be capacity shortages, and blackouts. Dependence on imports would raise costs. Deregulation is about to cause financial and technical problems.

In spite of all this, it seems that the sky has not fallen. This may therefore be a good time to review the situation, and see whether we should be concerned over the health of the utility industry.

Nationally, there is still a capacity problem, although it is not as severe as some observers feared. Areas with low reserve margins are localized. In some areas, load growth did not meet predicted levels, and in some areas new generation has been built. But there do remain regions that are not only short of capacity, they also lack the transmission facilities needed to import the energy required to solve the capacity problem (ECAR's Future, 1991). In the summer of 1990, there were voltage reductions and load shedding in the Baltimore-Washington area, because of capacity and transmission shortages at Baltimore Gas and Electric Company and Potomac Electric Power Company (1990 System Disturbances, 1991).

The utility industry is dependent on imports, both of energy and equipment. As long

as the imports take place in an environment of free trade, it is my opinion that this should not be a cause for concern. U.S. government involvement addressing these issues and supported only by arguments over restoring some sort of "all-American" utility industry are somewhat specious in the post-Communist "New World Order" of the late 20th century. (Arguments based on the impact on the nation's current-account balance may be more valid, but they take us beyond the scope of this report.)

Finally, there is the question of the financial health of the utilities. Perhaps because utilities operate in a largely regulated environment, there is little cause for concern here. A recent study by London Business School and *The Economist* (*The Best Companies*, 1991), compared 800 companies by added value as a percentage of sales. On this basis 4 of the top 10 and 12 of the top 30 companies were U.S. utilities.

Actually, the prospects for the industry are mixed. There is cause for some level of concern. The fact that the same Baltimore Gas and Electric that was mentioned earlier in connection with an east-coast power failure was in the list of companies with high ratios of added value to sales—in ninth place—must mean that financial health and reliable service are not perfectly correlated.¹⁰

1.4.1 Need for Government Involvement

The argument was made above that there is a need for additional R&D to be performed from the viewpoint of both the manufacturers and the utilities in this country. Let us turn now to a review of why government involvement is needed and justified.

First consider the situation of the manufacturers. Some U.S. companies have not been equal to the challenge of competition from abroad. This is particularly true in the heavy manufacturing side of the business. The power interests of Westinghouse, Square-D and many small companies have been bought by foreign concerns. In other companies the response has been appropriate: product lines have been rationalized, processes have been streamlined and product quality has been improved. Unfortunately, some of the infrastructure changes have come somewhat late, and R&D investments that should have been made over the last decade or so were not made. Lower than expected demand, and low demand predictions have not made the utility market attractive to long-term investment in the U.S. This is how the power part of Westinghouse came to be sold; and why the Swiss buyer (ABB—Asea Brown Boveri) is finding it necessary to pour millions of dollars into the venture to modernize it.

In terms of research, the continued involvement of the government is justified, it seems, by the evident lack of alternatives. EPRI, seemingly the logical alternative, is under pressure from the utilities to produce visible results. Consequently, they have tended to support short-term development. The kind of innovative, high-risk or long-term research that the government has sponsored in the past is just not being done elsewhere in the U.S. (But simply preventing foreign take-overs of a flagging industry should not be the proper role of the government. That some foreign buyers have "deeper pockets" than U.S. shareholders should not be a cause for concern—or surprise. And a viable Westinghouse that

¹⁰And it also supports Coase's viewpoint that the regulatory process has not worked to the benefit of the consumer.

happens to be owned by ABB will employ more Americans than a defunct Westinghouse. According to free-trade arguments, it is difficult to maintain that it is *essential* that the utility manufacturing industry in the U.S. be U.S.-owned.)

In terms of energy management systems and software, some U.S. companies have been acquired by European concerns (ESCA, SCI), but others are very strong internationally. Empros (a Division of Control Data Corp.), for example, has orders world-wide for their energy management systems. Power Technologies Inc. (a consultant) has recently opened an office in England, so as to be able to address the market of the European Community. Thus, it is very hard to argue that government nudging of any kind is needed here.

The same is not true of the utilities. Very likely because of the regulated environment they operate in, U.S. utilities have tended to become complacent. Planning strategies have not been optimized in terms of lowest societal cost, and have not considered conservation on an equal footing with system expansion. The idea of integrated resource planning as a tool for optimizing the use of the resources available to the utilities is very appropriate.

Because IRP is interdisciplinary in nature, and because it runs afoul of the organizational problems of most utilities, this kind of activity is an excellent candidate for government funding. The dollar invested here would be leveraged with the dollars spent by the utilities on conservation, and not spent by them on fossil fuel. In any event, there is a singular lack of alternatives for funding R&D in this area.

The same argument can be made for integrated operation. In a way, integrated resource planning and integrated operation represent solutions to the same problem, just with different time scales. Government funding of R&D for the integrated operation of the system should be just as highly leveraged as IRP funding. It is argued below that an effective way to accomplish technology transfer is to form R&D partnerships with industry, further leveraging the investment.

1.4.2 Technology Transfer

Suppose that some government-funded research group develops a new solution to some problem in the utility world. How best to get the technology adopted? Most researchers would want to write a paper. But research papers are written for other researchers, and few companies let their research groups decide policy. The paper has minimal commercial impact.

If we continue with the evolutionary model for R&D, the step that follows the new development is to test it in the real world. The problem is that utility conservatism will make this difficult. Truly it is said that the utilities want the very latest in high technology advancements, provided they've been tested on somebody else's system for 30 years.

The fear of the unknown will only be overcome by experience. The government-owned utilities—BPA, TVA and WAPA—are in a position to help, as they have in the past. BPA built one of the first Ultra High Voltage test lines, and TVA may still be the only U.S. utility with an optical CT. They should be encouraged to be innovative in this way. Given that not everything that comes from the research lab will be successful, its adoption by BPA or TVA should be seen as a test, and lack of success should not reflect badly on these utilities—either in terms of reputation or in dollars.

The ultimate object of demonstrating new technology is presumably to have it more

widely adopted. A company will add something to its product line. If too many reports are written and too many papers are published, no manufacturer can see how to turn the new technology into a product without all the competition being able to follow suit. So instead of the new ideas being tested in the marketplace, they never get beyond the demonstration phase. The implication of this is that if research work is to have any hope of commercial success, it may have to be kept confidential until it is brought to market or until it is protected by patent.

The desire for secrecy should be no surprise. If the research laboratory of a company had done the work, they would probably have been required to restrict the flow of information out of the company. The same approach may have to be applied to work done by the National Laboratories, for example.

Of course, the question of whether to write a paper or not might be moot if the work was done as a result of an R&D partnership already formed. Presumably in this event publishing rights would be controlled by the agreement with the company or companies that had funded the work. An advantage of the approach of forming a partnership before the work is done is that the necessary sanity check is performed on a continuous basis. The research is "directed." A difficulty is that the company is, to some extent, speculating on the outcome of the research. Conservative companies do not speculate.

It is encouraging to see that some new kinds of business relationships are developing between industry and government. The work on optical current measurement done at the Boulder laboratory of the National Institute for Standards and Technology (NIST), for example, has been developed exclusively by 3M Company. My group's work at JPL on a fiber optic communication system is also to be exclusively licensed.

It has taken an incredibly long time to find a buyer for the AbNET communication system developed by my group at JPL. There is a lesson to be learned from this, and I would like to generalize the result. Delay in having an invention adopted could be due to any of a number of causes:

- difficulty of finding the right buyer
- not-invented-here syndrome
- conservative nature of the industry
- the invention is not needed

It is inherently difficult to point to any one of these as the cause of the delay. The conservative nature of the industry has already been addressed. It would be presumptuous to address the not-invented-here problem.

If the invention is not needed, of course, it will be impossible to find the right buyer. But note that not all "needed" inventions will be recognized as such. (Xerography and lasers are two fine examples.)

The question of finding the right buyer may be solvable. At present, when the government wants it to be known that dollars are available for work in a certain area, a notice is inserted into a publication known as *Commerce Business Daily*, CBD. Companies eager to help the government spend their money read CBD, and respond if appropriate. What is needed is a publication that is a counterpart to CBD, that would be used to tell the private sector where the government was looking to them to co-sponsor some R&D. Let us call this new publication *Government Research Monthly*, GRM, for want of a better name.

It would take a while for this publication to be effective, to achieve the same reputation and readership as CBD. The proper promotional techniques would have to be developed. But it would be a worthwhile endeavor. A publication like GRM would add to the means available to begin the process of technology transfer, supplementing the other innovations being implemented at DOE, and relieving the researcher of some of the marketing chores.

This is not an argument against the publication of technical papers or reports. All of the traditional means of disseminating information are applicable to the kind of work that can be co-sponsored by government. But the timing should be considered carefully. Basic results could be published immediately in scientific journals. Applied results, especially those that have been co-funded by industry, should be published only after the commercial implications are understood.

The principal professional journal of the industry, the Transactions of the Power Engineering Society (PES) of IEEE, has a dual identity in this respect. It varies between preferring papers that represent a fundamental advance, and those of value to the practicing engineer. This makes it difficult, sometimes, to publish research results in a timely fashion.

1.5 Summary and Conclusions

Dramatic events on the world stage have recently shown the relative merit of the free market, the economic analog of the struggle for survival in biology. The power industry has to some extent managed to sidestep the evolutionary process, and finds itself struggling in the rapidly changing regulatory environment of the 1990s.

Regulatory changes proposed as part of the Department of Energy's National Energy Strategy can put the interconnected transmission system in the role of the marketplace for electrical energy. Fairly complete deregulation might be possible. This should have a stimulating effect on the industry, allowing competition of a kind that the utilities have hitherto not known.

This increased competition should motivate the industry to evolve again, to adopt some of the technical advances that they have, by and large, ignored in the past. But effective incorporation of new technology into a utility implies the availability of the technology in a form the utilities can use. Uncertain growth and lack of motivation have led to inertia, and some of the U.S. suppliers of the utility industry appear to be flagging.

Government support is justified not because the industry is facing foreign competition, but because an industry that is revitalized with government involvement can be cognizant of societal costs, while evolving in directions governed by market pressures.

Integrated power system control is an operational counterpart to integrated resource planning. Before the power system can be operated in an integrated fashion, a number of technical obstacles must be overcome. It is in the national interest to support the necessary work.

The size and complexity of the problem of operating the power system as an integrated whole will call for improvements in measurements, communications and computation. The problem can be broken up into separately solvable parts, involving optics, electronics and control theory. A systems approach can be used to integrate the results.

But it is inherently a hopeless task to *plan* the entire endeavor. I opened with a quotation about Darwin; I close with one from Adam Smith. The explanation of the

apparent movement towards harmony that Darwin exploited in explaining descent with modification is included in Smith's famous metaphor:

Every individual . . . by directing his industry in such a manner as its produce may be of greatest value, intends only his own gain, and he is in this as in many other cases, led by an invisible hand to promote an end which was no part of his intention . . . By pursuing his own interest he frequently promotes that of society more effectively than when he really intends to promote it.

The last sentence is the one I want to draw the reader's attention to. This section of the report is based on a paper I prepared for a DOE planning workshop held in November 1991. When the invitation from the organizers of this conference arrived at JPL, it asked that authors consider the power system of 2020, and where this industry was headed. I had difficulty in seeing this far into the future. Instead of having an overall idea of where we were going, I had several detailed ideas, candidates for a shotgun approach. The ideas needed a strategy for implementation.

The example of evolution has provided that strategy. It is not necessary, and probably not possible, that the entire system be planned at once. This kind of government involvement may lead to actions inimical to Adam Smith's invisible hand. However, if the involvement is to create options rather than to make choices, then market forces will still have the opportunity to work. The creation of options should be through government support for research and development. An evolutionary implementation approach should then lead to considerable excitement in the next decade or two!

SECTION 2

AbNET COMMUNICATIONS NETWORK

IN A GENERAL-PURPOSE distribution automation scheme, the need to communicate with all parts of the power system means that the topology of the power system fixes the topology of the communications network. The network can therefore be expected to include a large number of branch points, tap points and interconnections. These features make this communications network unlike any other.

The network operating software has to solve the problem of communicating to all the nodes of a very complex network, in as reliable a way as possible even if the network is damaged, and it has to do so with minimum transmission delays and at minimum cost. The solution adopted was a set of communications protocols that we called AbNET, described in an earlier report (Kirkham, Friend, Jackson, Johnston, 1990).

The details of the protocols were modified slightly during 1991. This section of the report will review the original design and describe the changes. In addition, two other advances in the network will be described. The I/O hardware to interface personal computers in the laboratory to a fiber optic network, and a software simulation of the protocols, both designed for the purposes of demonstrating the system, will be described.

2.1 The AbNET Protocols

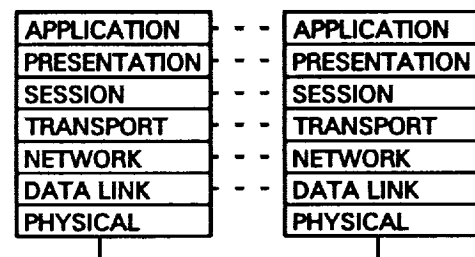
Fiber optics technology is well known in telephone trunk applications. However, the design of the AbNET network and the method of operating it are quite unlike the approach that the telephone company would use, because a different problem is solved. Trunks are point-to-point, the distribution system is not. The AbNET system was designed specifically for distribution automation. In the world of the phone company trunks, the design goals are to maximize the distance between repeaters, and to maximize the number of phone calls that can be squeezed into one fiber. This enables the phone company to hold down costs, and to spread the charges among a maximum number of revenue-earning calls. In distribution automation applications, AbNET uses repeaters wherever there is a need to tap or branch a fiber, so the distance between them is fixed by the power system and not by considerations of maximizing the use of the resource. The amount of traffic that the net-

work is expected to carry is low by the standards of fiber optics, so there is nothing to be gained by using a high data rate approach. Instead, AbNET maximizes the capability of the lowest-cost system. In terms of the revenue-earning efficiency of the phone company, AbNET is poor. However, its price and performance are appropriate to the application for which it was designed.

2.1.1 Open System Interconnection Model

The Open System Interconnection (OSI) model of the International Standards Organization (ISO) will be used here to provide a framework for discussion of the AbNET protocols. Computer-to-computer communications in the ISO seven layer model is shown in Figure 2-1. A more detailed explanation is given below. This description begins with the top layer and proceeds downward.

Figure 2-1
International Standards Organization model
of Open System Interconnection



- The **application** layer is the layer of software that needs to transfer data across the network. All other layers in the hierarchy exist solely to satisfy the needs of this layer. In the case of a distribution automation system, the applications layer will include the SCADA program, the program that presents data to the operator, and any control programs that may be used.
- The **presentation** layer provides services to the application layer to process the data in some way to make them more suitable for the layers below. This could mean, for example, translation or encryption of the data. (The operating system of the personal computer is an example of a presentation layer. It furnishes data in a format that can be handled, for example, by the word processing program or the screen driver, and at the same time can take a file of data furnished by a word processing program and present it for transmission to a remote computer to the layer underneath, the session layer.) In the distribution automation example, this layer could contain user-callable library routines for address translation.
- The **session** layer is the first of the layers in this hierarchy specifically concerned with communications to another computer. In essence, this layer is responsible for coordinating interaction between the opposite end application processes. The layer has to be aware of both the application and the communications. For example, a program sending blocks of data to make up a file might pause for acknowledgement at the end of each block. This could be a session layer function. In some applications, the session layer is a "virtual" layer. The decisions made here are the type of communication to be employed (e.g. full or half duplex) and how failures of lower layers in the hierarchy are to be handled.
- The **session** layer interfaces to the **transport** layer which is the highest of the seven layers responsible for the integrity of data. For example, this layer would be capable of

performing error checking, perhaps on a packet by packet basis. The job of the transport layer is to furnish error-free messages, in sequence, to the session layer. Whether this is a simple task or a complex one depends on the layers beneath. If the lower layers retain message sequence, and perform error checking, the transport layer becomes very simple. On the other hand, if the lower layers can operate so as to get messages out of sequence, as AbNET can for example, the transport layer must correct the deficiency.

- The transport layer connects to the **network layer**, which is responsible for furnishing data to the bare-bit manipulations of the data link layer below. The primary responsibility of the network layer is routing the information from computer to computer.

- The **data link layer** is the highest level at which information as such is handled. The data link layer may perform error control on a word by word basis; for example, parity checking occurs in the data link layer. There may also be some means of flow control or hand-shaking to assure synchronization between devices capable of operating at different speeds. The telephone company for example, in their data link layer, typically will add bits to the 8-bit word coming from most modems, so that error checking (in addition to the parity bit that the user knows about) is usually performed.

Some users of the ISO model have divided the data link layer into two sub-layers: **logical link control** and a **medium access control**. The logical link sub-layer is responsible for establishing, maintaining and terminating a logical connection between devices. In a parallel connected network, the medium access control sub-layer ensures that only one device attempts to transmit at a time, performing the function of congestion control.

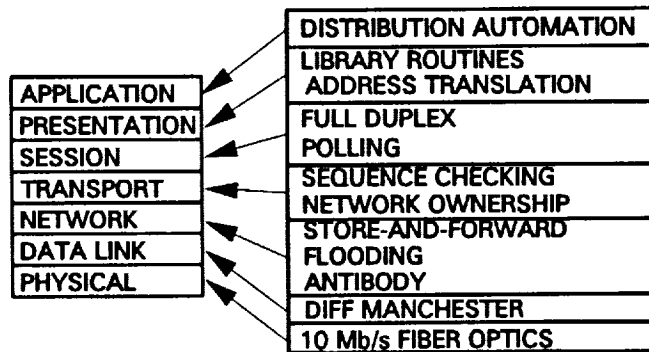
- The **physical layer**, the lowest layer of this hierarchy, is the level at which the electrical or optical signals are exchanged. A specification of the physical layer typically includes a description of electrical, optical and mechanical quantities involved. For example, RS-232 describes the type of plug to be used as well as the voltage levels and the pin connections. In the world of modems, the physical layer includes a description of the procedures to establish or release connections between electrical circuits such as phone lines.

One may note that the functional relationships between layers are clearly defined in the open system interconnection model, but are often not so clearly defined in practice. The visibility of the boundaries between layers need be clearly defined only if the boundary corresponds to a product, in which case interfacing to that product (either physically or in terms of software) will be vastly simplified. Thus, the transport layer and the session layer functions are both performed adequately by the communications programs used in personal computers; and the network layer, data link layer, and physical layer occur outside the boundaries of the typical modem.

2.1.2 Review of Original AbNET Design

AbNET is a dual-hybrid network, comprised of a fiber-optics-and-power-line-carrier digital data system in which every node (or RTU) is also a repeater. A passive fiber optics PCM (pulse code modulated) voice channel shares the fiber cable. The nodes are interconnected by fiber optic links routed along the distribution system. The fixed installation of the voice channel is intended to be passive, the optoelectronics and the electronics being carried by the user and inserted into the system as needed. Only the digital data system can be put in the framework of the ISO OSI model. This is shown in Figure 2-2.

Figure 2-2
Main features of AbNET
in OSI framework



- Distribution automation is shown in Figure 2-2 as the applications layer. The bulk of the routine traffic will probably be associated with data acquisition. Descriptions of the various other functions that comprise distribution automation need not be repeated here. Suffice it to say that AbNET was developed explicitly to satisfy the communications requirements of data acquisition and control for distribution automation. The brief descriptions that follow show how AbNET meets these requirements, and supports distribution automation.

- The principal function of the presentation layer is to perform the conversion of data descriptions (in application programs) into addresses that the lower layers can recognize.

- The session layer establishes transport connections between the central unit and the remote nodes on a polling basis. Centralized information about network performance can be made available to an applications program only through the operation of this layer. The session layer therefore retains the ability to establish (and test) connections between adjacent nodes. Since message delivery can occur in the AbNET system even if a number of connections are broken, network integrity is checked by software, operating transparently to the user, at this level.

- The transport layer is centralized at the distribution substation. It normally operates through the local network in support of the data acquisition and control functions. Under unusual circumstances, it can choose to use the local network for access to the area served by another substation.

Normally, part of the packet header is the node address, a 2-byte logical designation. One of the two bytes is used in AbNET to indicate "ownership." A substation owns all the RTUs in its service territory, and none belonging to the neighboring territory. Ordinarily, network layer operation screens the higher layers from messages intended for RTUs in other areas. Along the border of two service areas, however, the RTUs (called gateways) have to respond to two (or more) master stations at the substations. This means that the rejection of "foreign" messages occurs not at the boundary itself, but one layer of RTUs deep into each service area.

This feature provides a mechanism for one master station to take over the territory of another if it is ever needed. A special kind of message can be sent from one master to the RTUs of another, re-designating ownership, and redefining the boundary.

- The network layer is a decentralized message-based, store-and-forward system. Congestion in the network is controlled by recognition of ownership, and by an antibodylike algorithm, that allows a node to repeat a message once and only once. Routing is accomplished by network flooding.

- The data link layer incorporates the field descriptions and framing definitions of

IEEE 802¹¹ as far as possible. Because the bit error rate is expected to be extremely small, there is no handshaking for flow control, and no request for retransmission in the event of an error. This avoids system problems in the case of hardware failures.


- The physical layer is a fiber optics based hybrid. Multimode fiber can be used, transmitting optics will be based on LEDs operating in the near infra-red, the bit rate will be 10 MHz, and differential Manchester coding will be used. (In many respects this is similar to Ethernet). Access to locations on the power line secondary will be by means of low-power high-frequency power line carrier communications.

2.1.3 Additions to AbNET Design

In operation, the application program performs its function by passing instructions through the presentation layer to the session layer, for coordination with the communication system. The session layer is responsible for polling the RTUs, and gathering data from them. Although it was never stated explicitly, it has been generally assumed that the polling would be sequential, with each RTU being addressed in turn. This would follow naturally from a program that expected data from each RTU in turn. While this will work, and with the AbNET physical layer it is fast enough to meet the requirements of SCADA, a speed improvement is possible at the cost of a slight decrease in reliability. This modification permits the application layer to issue an instruction equivalent to "Poll all nodes" to the lower layers.

Consider the section of a network shown in Figure 2-3. A master station has control of some RTUs that happen to be connected in a loop. If each is addressed in turn, the polling sequence would be: Poll 1 (wait for response); poll 2 (wait for response); poll 3 . . . and so on.

Figure 2-3
Example of simple network



The message handling sequence is more complex, however. Thus, the sequence begins: Master to 1; 1 to Master; Master to 1, 1 to 2; 2 to 1, 1 to Master; just to poll the first two RTUs. As the length of a line (in terms of the number of RTUs on it) increases,

¹¹For most purposes, and in the case of most standards, it is only the bottom three layers of this hierarchy that are standardized. However, because these layers do need to be standardized for effective communication between different machines, a good deal of work has gone on in this area. IEEE in particular has developed a series of standards according to IEEE Computer Society Project 802. These are usually known by the decimal organization of the subcommittees. For example, Ethernet has been adopted as IEEE 802.3 or CSMA/CD (Carrier Sense Multiple Access with Collision Detection) and IEEE 802.5 is an embodiment of IBM's Token Ring.

the number of "hops" required to poll it increases faster. Considering a linear network, the number of hops required to poll the RTUs is as shown in Table 2-1.

Table 2-1
Number of hops required to poll a linear network

nodes	hops
1	2
2	6
3	12
4	20
5	30
6	42

In a linear network, such as the one shown in Figure 2-3 if the connection from the master to node 8 is broken, the distance from the master to node j is j lines. To poll this node, and to obtain the response, therefore takes $2j$ hops. The number of hops H required to poll all the nodes up to node n is therefore simply

$$H = 2 \sum_{j=1}^n j$$

It can be shown¹² that this expression reduces to $n(n+1)$, which greatly simplifies evaluation. In the network of Figure 2-3 the most distant node is only 4 lines from the master. The network will therefore be polled in 20 hops. In a system that has 10 lines between the master station and the most distant node, it will take 110 hops for the last response

¹²By definition

$$\sum_{j=1}^n j = 1 + 2 + 3 + \dots + (n-2) + (n-1) + n$$

reversing the order:

$$\sum_{j=1}^n j = n + (n-1) + (n-2) + \dots + 3 + 2 + 1$$

adding term by term:

$$2 \sum_{j=1}^n j = (n+1) + (n+1) + (n+1) + \dots + (n+1) + (n+1) + (n+1)$$

There are n terms in the series, so that:

$$2 \sum_{j=1}^n j = n(n+1)$$

to be received.¹³

The total number of hops for a data acquisition scan is reduced if the flooding algorithm of the network layer is modified. In the general flooding network, a node acting as a repeater retransmits an incoming message on all outgoing lines. Similarly, it transmits its response to a message addressed to it on all outbound lines. In the network of Figure 2-3, for example, node 2 responds to the master on both the line to node 1 and the line to node 3. Unless there is a network failure, only the response to node 1 is important.

Taking advantage of this, suppose the action at node 2 consists of a response to the master, on the line to node 1, and a simultaneous repeat of the poll on the line to node 3. This change has the effect of reducing the number of hops required to poll the network to the number required to poll the most distant node. The system of Figure 2-3 can be polled in 8 hops; a 10-line system can be polled in 20 hops. The improvement due to the modified flooding algorithm for these two cases is thus between a factor of 2 and a factor of 5.

There is a drawback, but it is relatively minor. If the network configuration changes between the reception of the poll and the end of the response, messages could be lost. Suppose the line between the master and node 1 fails immediately after node 1 receives the request for data. In this case, the responses from nodes 1, 2, 3 etc never arrive at the master. Redundant copies are not circulated in the other direction. While a copy of the poll might arrive at the low-numbered nodes from the other direction, these nodes will judge that they have already responded. However, within a few ms the master can detect that node 1 has "timed-out," and that several of the nodes have not responded. If a new poll is generated (and if the network does not change again), the missing data can be obtained, the poll being routed automatically by the flooding algorithm.

The loss in reliability is relatively minor because data are lost only when the network configuration changes. Once the configuration has stabilized, the two versions of flooding behave identically. Field testing of the two versions may indicate the relative value of the performance improvement and the reliability loss.

Another, related, concept worth considering in the context of squeezing the maximum performance out of the AbNET protocols is "rapid-fire" polling. In a full-duplex system, there is no need for the master station to wait for a response before transmitting a new poll. One can visualize a network carrying an endless train of polling messages away from the master, and a more or less continuous stream of data returning in reply. In the original polling scheme, the master station waited for each node to respond before the next was polled. In the revised polling scheme, the master station waited for a response from the set of all nodes before starting a new scan. It is difficult to estimate the speed improvement that would result from waiting only for the output channel to become clear before sending out a new poll. An assessment may be made in the near future using the software simulator (described below).

One further refinement is worth mentioning. In a practical distribution automation scheme, there are likely to be many RTUs dedicated to customer interaction, such as meter reading or load management. These RTUs are nodes that can be expected to be connected in

¹³Even this large number does not result in unacceptable delays. If each individual message is 1000 bits long, it takes 100 μ s to insert the message into the network. With no allowance for processing, 110 hops thus takes 11 ms.


a linear fashion along the feeders, contributing significantly to the number of nodes to be polled in a single scan. In terms of system operation, such RTUs will likely have little data of value. The AbNET protocols therefore recognize different classes or priorities of RTU. During normal SCADA operation, the low-priority nodes will act as repeaters, adding slightly to the message transit time. They will not furnish data unless they are specifically addressed.

In the original flooding scheme, not polling low-priority nodes could lead to significant reductions in the scan time of the system. In the modified protocol, there would be no effect other than a reduction in the amount of data delivered to the master.

2.2 Software Simulator


In order to make a convincing demonstration of the operation of the network protocols, a representative topology was planned. The network should have a large number of intersecting loops or meshes, and a small number of nodes that were served by spurs. Since the network was supposed to represent part of the distribution network served by a substation, it should have more than one line from the master station. While the underlying power system could be assumed to be operated radially, the communication system could bridge open disconnects. In collaboration with George Allen of American Electric Power Service Corporation,¹⁴ the configuration shown in Figure 2-4 was decided upon.

Figure 2-4
Representative portion of
communication network



In order to show the flexibility of the network, a variation with two master stations was also designed. This is shown in Figure 2-5.

Figure 2-5
2-Master version of
communication network



¹⁴At the time, AEP was considering a "real-world" demonstration of the AbNET network, to be held at their John Dolan Laboratory near Columbus, Ohio.

In this version, nodes numbered 2, 7, 12, 13, 18 and 23 are of the type we have designated "gateways." In the AbNET system, a gateway differs from an ordinary node in that it responds to two (or more) masters. In the Ethernet terminology, it is "promiscuous." Gateway nodes are located along the border between two service territories. Rather than put two RTUs on a pole at a border, each owned by one particular substation, the gateway RTU is owned by both substations. Not only does this make economic and aesthetic sense, it adds the flexibility to move the boundary by means of commands sent over the network itself. There are a number of quite realistic scenarios in which such a feature could be advantageous in distribution automation.

2.2.1 Overview

The original purpose of the software simulation (as opposed to the demonstration in Ohio) of this network was to investigate the impact of operating a half-duplex protocol. While it had been intended to operate full-duplex from the very beginning, the laboratory scale demonstration, using modified Ethernet hardware and ordinary PCs, could be done more easily and at lower cost if the protocol were half duplex. The question to be answered was How much impact would half-duplex operation have on system performance. Accordingly, a simulation was written in which the network of Figure 2-4 and the networks of Figure 2-5 were modeled in considerable detail.

The simulation was written in QuickBASIC, for convenience and portability.¹⁵ The three major parts of the program will be described.

- First, the initialization section, executing once, is used to determine which network to simulate. There are a number of other user options. Since the simulation uses only one processor to emulate the behavior of the 25 processors in each network, it was necessary to examine each in turn, to check for messages, generate responses and so on. Because this meant a deterministic approach to the performance of the system, the option was included to have the master stations generate destination addresses either in sequence, or at random. The option is also offered to take lines out of service, or restore lines to service that have been removed. After initialization, a message is created and stored at the master (or masters), and a flag is set to indicate to the next part of the program that a transmission is ready.

- The remainder of the code operates iteratively. The simulation portion begins by examining each node to see if it has a message to send. Flags have been set (either by the initialization section, or by the previous iteration) for this purpose. If a transmit flag is set for a node, the message it has to send, stored in a transmit buffer, is copied to the connected nodes, unless a flag derived from considerations of full or half duplex blocks the copying. This is done for all nodes.

Next, the program scans each node again to determine what action it would take operating under the control of the AbNET protocols. For example, if a node has received

¹⁵At the time of writing the simulation consists of about 3000 lines of code. The complete source code will not be published in this report. Readers who are interested in experimenting with the simulation, may obtain a copy by sending a DOS-compatible disk to the author, within a year of the publication date on the cover.

a message, it checks to see whether it should respond, repeat the message, or ignore it. At this point, a new set of transmit flags is set, and the iteration repeats.

- As the simulation executes, a display portion is called to show the operation of the protocols. A map of the network is shown on the screen, with the nodes numbered, and with downed lines indicated. Each message is associated with a color, and the color corresponding to the last message handled at any node is used to fill the box representing the node. In addition, a vertical bar showing the source and destination of the last eight messages, along with their color, is displayed for each master station. A tally is kept of the number of nodes that have responded to the poll.

2.2.2 Results

The original purpose of the simulation was to see whether the performance of the system would be adversely affected by half-duplex operation. At the time of writing, only the original AbNET protocols (the ones designed prior to the modifications described in Section 2.1.3) have been simulated.

One of the first results of the simulation was that, in the single-master network, exactly the same number of nodes could be scanned in a given time in the half-duplex version as the full-duplex version. In other words, there were never any delays in transmitting messages because of having to wait for traffic in the other direction.

This may at first seem surprising. The result arises from the combination of polling and flooding. This combination means that no RTU can spontaneously originate a message. Messages flood out from the master, and when a copy reaches its destination, the response floods back. There is a possibility of two messages arriving at an RTU at the same time, but they will arrive over different fibers. This is permitted, and does not result in a loss of data.

There is a small defect in the half-duplex version of the model that should be mentioned. Because each node is examined in turn, low-numbered nodes are allowed to send, and the higher-numbered nodes they send to are made to wait. This is somewhat unrealistic. In a real network, there would arise some occasions when the two nodes attempted to send to one another at the same time. In the absence of handshaking, neither message would arrive.

In the example of Figure 2-3, nodes 4 and 5 would attempt to send a copy of the polling message to each other at the same time. While this would mean that neither message arrived, it is also clear that there is no loss of information, since both nodes have already handled the poll in any case. It is generally true that a collision of this sort, concerning identical data, will not result in the loss of information. It can only occur over links that happen to connect two parallel paths.

On the other hand, an overlap on a link could happen in the case of multiple masters, where the messages in the area of the gateways are different, but are ready to be sent at the same time. Consequently, it can be argued that the results of the model are optimistic in the case of two masters, since the code allows messages to get through that should not. In defense of the simulation, it can be said that it contains a compensating bias towards pessimistic results in the two-master case. In the model, an unrealistically large number of the RTUs are gateways. The model gateways are therefore obliged to handle

a larger fraction of the network's total traffic than is realistic. Because the program examines low-numbered nodes first, high-numbered nodes can be made to wait. The effect of message overlaps is therefore exaggerated.

The model shows that half-duplex operation is not more than 10% slower than full-duplex, for the 2-master network of Figure 2-5. Since this network has the high proportion of 6 gateways for 17 ordinary nodes, it is reasonable to expect that a realistic network would perform rather better.

Of greater interest to the operator is the demonstration by the simulation of the automatic routing feature of the protocols. Arbitrarily chosen lines can be taken out of service, and new routes to the destination are found. With colors representing messages, this becomes a fascinating exercise.

2.3 I/O Adapter

In the last of these reports (Kirkham, Götsche, Niebur, Friend and Johnston, 1991), the similarity of the lower layers of the AbNET protocols to Ethernet was noted. Taking advantage of this, it was proposed to demonstrate the AbNET protocols by means of a network of PCs, connected by fiber optics. At each PC, the fiber cable would be interfaced to the PC by means of a modified Ethernet board, and an I/O adapter of JPL design.


A number of I/O adapters have been made and tested. The circuit of the final design is shown in Figure 2-6.

Power is taken from the 15-pin connector on the Ethernet board. It supplies two essentially independent circuits, for the receive and transmit sides of the interface. Both circuits are fairly conventional, and require little comment here. One unusual feature can be seen in the receive circuit. The optical/electrical converter is an integrated device, but it does not provide sufficient output to drive the CMOS circuits required to interface to the computer. A circuit of a type known as MMIC (monolithic microwave integrated circuit) has been used to provide gain.


These circuits were commercially developed as RF amplifiers. They provide moderate amounts of gain (usually <20 dB) very conveniently. They are typically 4-terminal devices, with 50- Ω inputs and outputs. They make very useful and simple broadband gain blocks, with constant gain up to or even above 1 GHz.

In this application, their bandwidth makes it possible to maintain the steep wavefronts of the 10 MHz signals on the fiber. It is hard to think of a simpler and more stable means of amplifying the small signal coming from the receiving diode to the level needed by the remainder of the circuit.

(a) Power supply



(b) Transmitter



(c) Receiver

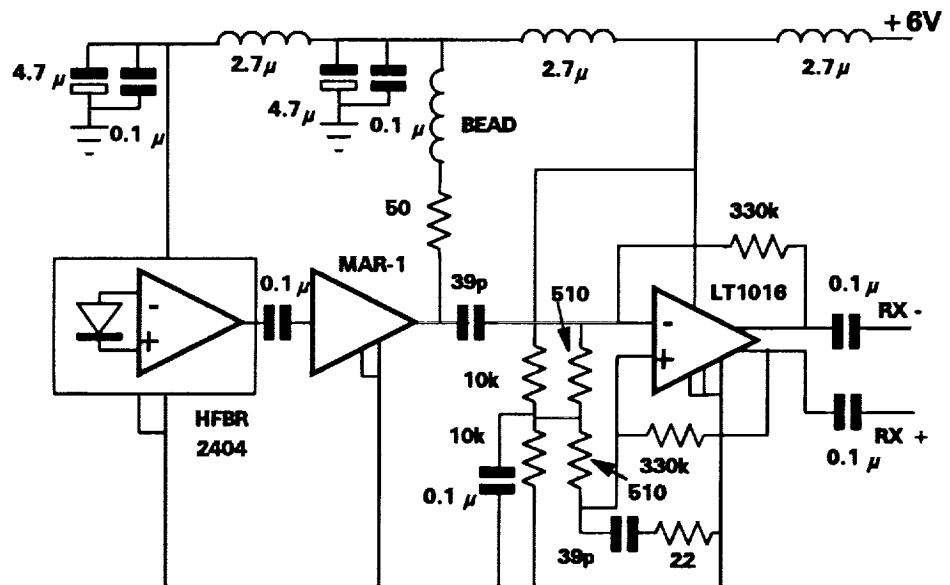


Figure 2-6 Circuit of fiber I/O adapter

SECTION 3

INTERFACING MEASUREMENT TRANSDUCERS

SOMEWHERE about the turn of the century, instrument transformers began to be used in electric power systems. In those days, the power system operator controlled the system by looking at instruments on a marble display board, custom built for the particular power system. The meters were brass, and the full voltage and current of the power system were applied to the board. *Power system operator* became a hazardous profession as system voltages were increased—until instrument transformers became widespread.

Current and voltage transformers were introduced to permit low-voltage meters to be used with relatively little insulation. Casual contact with the display panel became considerably safer!

Two standards emerged. Current transformers were standardized to furnish 5 A into their secondary burden at rated primary current, and voltage transformers were standardized at 120 V at rated voltage. I have not been able to ascertain when these values were first used, but by 1911 the National Bureau of Standards in the U.S. was writing internal reports (Agnew, 1911) discussing error sources on instrument transformers with these ratings.

To put the period into perspective, remember that at the time that these standard values were emerging, attempts were being made to record and play back sound. These early efforts were entirely acoustical in nature: the sound of the artist's voice, operating on a small stylus, caused a wavy line to be engraved in a wax cylinder. When another stylus, attached to a large horn, retraced the line, sound could be heard.

Later, disks rotating at 78 rpm replaced the cylinders. While these disks were entirely incompatible with the cylinders, the basic technology was the same. The disks produced better sound, and were easier to store.

Vinyl disks, rotating at 45 or 33 rpm, supplanted the "old 78s." Electronic amplifiers replaced the horns. New equipment was required to play the recordings.

At each step, the older standard was not replaced, it was supplemented. However, the quality was so improved that eventually the older standards were abandoned. We may suppose that the new technology of optical compact discs will eventually replace long-playing records.

While these advances in the reproduction of sound were going on, the electric supply business was frozen into its original standard. 80 years later, it still is! Was there such a

large investment in the old technology that it was unthinkable to make it obsolete? Probably not. The industry's size doubled every ten years for most of this century, so that ten years after any new standard came into effect, half the system would be using it. Is the industry so very conservative? Yes, see Section 1. But the development of optical current transducers over the last few years by a number of manufacturers has given rise to the need to reexamine the question of interfacing interface "transformers." It seems certain that the earlier 5-A standard will not be retained, so the question is, Is a new standard needed to replace the old one, and if so, what? This section of the report will examine the issues, and show how a decision could be made.

3.1 Why not retain the 5-A Standard?

The use of a 5-A current transformer was justified principally by two factors. The first was *noise immunity*. The level of current was sufficiently high that a false signal of similar magnitude could scarcely be induced in the pilot wires. In protection and metering, it was not difficult for the equipment to distinguish between a 5-A signal and a smaller noise that might be coupled into the circuit. While "primary relays" could be used in the low voltage part of the system, relays and meters driven by instrument transformers were more practical at high voltages. With a lower current, the relatively long pilot wires in a differential protection scheme, for example, might be a source of false trips.

Second, the *energy* available from the instrument transformers could be used to actuate both the metering and the protective relaying equipment. "Signals" at 120 V or 5 A can be used to create quite powerful electromagnets. A range of ingenious electromagnetic devices has been developed to meter and protect the various parts of the power system, all based on high-energy signals from standardized instrument transformers.

There seems to be growing understanding that optical or electronic measurements should make no attempt to reconstruct the high-power 5-amp or 120-V signals that would enable them to meet the existing standards.

- First, there is no need to. The required noise immunity is provided by the nature of the optical fiber itself. As a carrier of information, and particularly digital information, the few microwatts of optical energy in the fiber are remarkably immune to external influence. And the energy required to operate a modern electronic relay is extremely small compared to that needed for an electromagnetic device. (Note that while some modern solid-state relays are built for 5-A, 120-V operation, they contain transformers that convert the input signals to low-level values for internal use.)

- Second, it is scarcely practical to produce a 5-A replica. For a single current transducer, there would be no difficulty, but for the large number required in a typical substation application, the power supply design would be, to say the least, challenging. Since, in the case of CTs for protection, the current transformer output is supposed to be a replica of the input even if the input is 20 times larger than its nominal value, the power supply has to be able to deliver at least 100 A per CT. And even assuming that a power supply could be made, the impact on the station battery of having several such devices connected would be detrimental. Thus, the community of users of instrument transformers has concluded, informally at present, that a new interface will soon be needed.

This being the case, the question naturally arises as to what standard should be met by an optical or electronic measurement. Is a new standard needed at all?

3.2 The Need for Standards

In practice there are several kinds of standards, instituted for a variety of reasons. Some of them are listed below.

- **Continuity standards** are a simple way of ensuring that the same product is used over the years. For example, it is unimportant what the refractive index of transformer oil is, but it may be specified to ensure continuity of supply. This is a substitute for a
- **Performance standard**, that would ensure that the right material was used for the job in question by examining its ability to perform in that application. The fairness of standards of comparison is ensured by
- **Weights and Measures standards**, that are written to see that a volt is the same in Denmark as it is in the U.S., and so on.
- **De facto standards** come into being as the result of market forces. A technical specification that becomes widely known and understood becomes a *de facto* standard. This is likely to happen if there is one major supplier of a market, or one dominant user. Such standards often become the subject of
- **Enshrining standards**, which take the work of others, usually in the form of *de facto* standards, and give the work the stamp of approval. IEEE standard 488 is such a standard, derived from Hewlett-Packard's HP-IB interface specification.
- **Enabling standards** spell out processes that have yet to become popular enough to become *de facto* standards. The decision by the FCC in the U.S. to allow color TV transmission according to the NTSC standard was an enabling standard.

There are a few other kinds of standards. **Safety standards** are written to protect people from dangerous environments, and **environmental standards** are written to protect the environment from dangerous people. Standards can be used to support industrial protectionism or international science. For an area that has such wide-ranging implications, standards get surprisingly little attention.

There is no *de facto* standard for optical current transducers, perhaps because there is no single, dominant supplier or user. Users would feel more comfortable if they knew their transducers were built to an interface standard that gave them some flexibility of application. Manufacturers would be happy not to have to custom build output circuits. It is thus the consensus¹⁶ of the users and the manufacturers that a standard is needed for optical current transducers. Is the timing right for an enabling standard? Or if we wait a little longer, will a standard emerge?

3.3 Timing of Standards

The timing for enabling standards frequently involves a compromise between setting the standard too early, so that the industry is unable to take advantage of technological developments, and setting the standard too late, so that the market becomes fragmented, with each manufacturer having his own standard. Early standards may limit innovation, late standards may hinder market development.

¹⁶Obtained, for example, at meetings of the Optical Sensors Manufacturers and Users Group (OPSMUG), organized by Emile Hyman of Public Service Gas and Electric in Newark, NJ.

In practice, standards in a developing industry leap-frog one another. There are examples in the consumer electronics industry. European color television is better than American television because it uses a standard (PAL) that was developed later than the American NTSC. The technology had advanced between the times that the two standards were adopted. On the other hand, America had color television for several years before it was available in Europe. In a similar vein, Britain had black and white television according to a 405-line standard established before WW II, which proved later to be incompatible with color. When Britain wanted to provide a color service, an entirely new set of standards had to be developed. The 625-line PAL standard was adopted.

The standards for instrument transformers in the power industry have remained unaltered for 80 years! The accuracy requirements for metering grade transformers, and the dynamic range requirements for relaying transformers may have become more stringent during the period, but the basic concept of 5 A for CTs and 120 V for PTs at rated value has not changed since shortly after the turn of the century. This situation is very unusual. Indeed, it is without parallel, because electrotechnology has advanced considerably during the same period.

3.3.1 The Role of the Market

It is tempting to think that "the market" can be left to decide what the standard should be for low-energy measurement transducer interfaces. This is unlikely. Consider the interfaces that have been used in four efforts (Berkebile, 1980; Houston, 1980; Deliyannides *et al.*, 1981; and Schweitzer *et al.*, 1981), funded by just one entity, the Electric Power Research Institute (EPRI). The interfaces are described in Table 3-1.

Table 3-1 Interfaces used in selected EPRI projects

Project or Report	Interfaces Used	
	Transducer to Receiver	Receiver to Utility
RP 560-1 (EL-1611)	12-bit serial with "range" bit, 1 MHz clock rate	12-bit parallel digital (15-V CMOS) and ± 10 -V analog on two ranges
RP 668-1 (EL-1343)	14-bit serial, 8-kHz sampling rate	± 20 -V analog, and 14-bit (5-V TTL) parallel digital
RP 1359-1 (EL-1813)	not addressed	16 bits (for revenue). Also asks for an interface to an RP-560 electronic CT, and analog high current (3-A and 30-A) tripping outputs.
RP 1359-3 (EL-1601)	12 bit 240 or 480 samples per second	one output number per second on LED display

Dr. Narain Hingorani, Director of EPRI's Electric Systems Division, told the author a few years ago that "Standards follow practice," and that it was not EPRI's role to set standards. It seems clear that neither EPRI nor market forces are going to decide on an interface standard. The industry does not have a single dominant supplier (though ABB may be close to that in some areas) nor a single dominant user. This would certainly argue in favor of an enabling standard.

This is a subject that is very timely. The first OCT device on the market was from ABB. Now it is only one of several optical measurements. In the U.S., Westinghouse (now also part of ABB), Square-D (now part of Schneider) and 3M have all demonstrated optical CTs (Cease and Johnston, 1990; Ulmer, 1990, MacDougall, Lutz and Wandmacher, 1991). There are also several European and Japanese companies with products on the market. Utilities, always conservative, feel uncomfortable enough with the new optical sensing technology. The lack of a defined interface standard makes matters worse. It is likely that the market for electronic or optical measurements will be retarded by the non-availability of interface standards that people can build to. This argues in favor of an enabling standard in the near future.

3.4 A Close Look at the Existing Standards

ANSI/IEEE C57.13-1978 (American National Standards Institute, 1978), is a specification for *instrument transformers*, not just for an interface. Further, it deals with both CTs and PTs, and covers relaying service as well as metering service. The International Electrotechnical Commission has three relevant standards: IEC 185: 1987, Current Transformers; IEC 186: 1987, Voltage Transformers; and IEC 44-3: 1980, Instrument Transformers, Combined transformers. These, too, are specifications for the performance of the instrument transformer, rather than just the interface.

Since it may be effective to model a new standard after one of these, let us look at each one, to see what it requires. In ANSI C57.13 there are 9 major sections, as follows:

1. Scope
2. Definitions
3. References
4. General Requirements
5. Accuracy Classes for Metering Service
6. Current Transformers
7. Voltage Transformers
8. Test Code
9. Bibliography

The titles of the first 3 sections—Scope, Definitions and References—speak for themselves. The General Requirements have to do with air density, altitude, temperature rise and insulation; so for our present purposes of discussing performance requirements, we can concentrate on sections 5 (Accuracy Classes for Metering Service), 6 (Current Transformers) and 7 (Voltage Transformers).

IEC Publication 44-3 (combined transformers) deals with requirements and tests that

are additional to the requirements of IEC 185 and 186 that are necessary for a transformer consisting of a voltage transformer and a current transformer in one case. Therefore, let us concentrate for now on IEC 185, current transformers. This document has 15 major sections, divided into three Chapters, as follows:

Chapter I: General Requirements Applicable to all Current Transformers

1. General
2. Rating and Performance Requirements Applicable to all CTs
3. Tests, General
4. Type Tests
5. Routine Tests
6. Special Tests
7. Marking

Chapter II: Additional Requirements for Measuring Current Transformers

8. General
9. Accuracy Requirements
10. Tests for Accuracy
11. Marking

Chapter III: Additional Requirements for Protective Current Transformers

12. General
13. Accuracy Requirements
14. Tests for Accuracy
15. Marking

Broadly, Chapter I corresponds to the first four sections of the ANSI standard, so that we can concentrate on Chapters II and III, which deal separately with CTs for metering and relaying.

Despite the extensive amount of space devoted to other topics, the essence of these standards is, of course, the requirements in terms of accuracy. ANSI defines a number of *correction factors* to help understand the requirements. The *phase angle* correction factor is the ratio of the true power factor to the measured power factor. This is described as being "a function of both the phase angles of the instrument transformer and the power factor of the primary circuit being measured." Seemingly of more practical use is the *ratio* correction factor (RCF). This is the ratio of the true ratio to the marked ratio. By way of explanation, ANSI adds that "the primary current is equal to the secondary current multiplied by the marked ratio times the correction factor." The ANSI definition can be written

$$I_p = (I_s K_n) \times RCF \quad (3-1)$$

where K_n is the marked transformation ratio, I_p is the actual primary current and I_s the actual secondary current when I_p is flowing. Thus

$$RCF = K/K_n \quad (3-2)$$

where K is the actual transformer ratio.

The IEC document defines a *current error* or *ratio error* by the formula:

$$\text{Current error } \% = \frac{(K_n I_s - I_p) \times 100}{I_p} \quad (3-3)$$

where K_n , I_p and I_s are as above. The difference between the two treatments is that ANSI treats the currents as being known, and derives a “fudge factor” for the transformer ratio so that the true primary current can be calculated, whereas IEC treats the transformer ratio as fixed, and says that the secondary current contains an error in its estimate of the primary current. Since ANSI uses the ratio correction factor to express the limits of accuracy, and IEC uses the limits of current error, we need to relate the two terms.


If the actual transformation ratio K is, say, high by 1%, for a 500:5 CT, the ratio would actually be 505:5. Equation 3-2 shows that RCF would have to be 1.01. The estimate of the primary current would be low by a similar amount. In other words, with 5 A in the secondary, one would expect there to be 500 A in the primary. In fact, there would be 505 A. The IEC error % is related¹⁷ to the ANSI RCF by

$$\text{error } \% = (RCF - 1) \times 100 \quad (3-4)$$

For metering service, ANSI defines 3 accuracy classes: 0.3, 0.6 and 1.2 for which the maximum values of RCF are $1 \pm 0.3\%$, $1 \pm 0.6\%$ and $1 \pm 1.2\%$ at rated current. Twice the correction is allowed at 10% rated current. The permitted phase error is related to the RCF by means of a diagram.

IEC defines 6 accuracy classes: 0.1, 0.2, 0.5, 1, 3 and 5. These numbers are the limits of error (in percent) at rated current. The limits of error for other current values are specified in a table. Some of the values are combined with the ANSI diagram in Figure 3-1.

Figure 3-1
Accuracy requirements
for metering service




¹⁷This expression is not in agreement with a footnote in the ANSI standard, which states that “The error of a transformer, rather than the correction, appears in much of the literature. It is the negative of the correction and is ascribed to the ratio I_2/I_1 .” It seems that the terms must be different: if the RCF is close to unity, the error is close to zero.

Direct comparison of the limits is not possible, since the ANSI and IEC accuracy classes at 100% rated current are not identical, and ANSI specifies the performance at 10% whereas IEC specifies both 5% and 20%. However, Figure 3-1 shows that the requirements are broadly similar.

An alternative presentation of the accuracy requirements is shown in Figure 3-2. In the figure, the permitted error is shown as a function of current. The points at which the requirements are specified have been joined (somewhat arbitrarily) by straight lines, even though a logarithmic axis has been used for the current.

Figure 3-2
Accuracy requirements
for metering service
as a function of current



The general trend of Figure 3-2 is clear. At rated current a certain error is permitted, and at lower currents the standards are more tolerant. The allowed error is specified only at a few points because the CT can be presumed to be "well-behaved" and not subject to abrupt nonlinearities. There has to be some allowance for nonlinearity in the devices, which are based on magnetic circuits with saturable iron cores.

While these accuracy limits may be applicable to a low-energy sensor, a different approach may be warranted. With an electronic output, for example, the linearity can be expected to be excellent. The accuracy at rated current can be viewed as a matter of adjusting the gain. (This could be as simple as turning a control.) If there is also an offset control, the error can be made arbitrarily small at two values of current.

By thinking of the allowed error in this way, it becomes clear that Figure 3-2 is simply a specification of the inaccuracy permitted in the *scale factor*. IEC standard values of secondary current are 1 A, 2 A and 5 A, with 5 A being the preferred value. ANSI is rather more vague: 5 A secondary current is assumed (but not made a requirement) and guidance on how to adapt the standard to other currents is given.

Thus, to return to the central question of this section of the report, if we replace the secondary current of 5 A with some other value (such as 5 mA) or even some other parameter (such as 5 V), what should the choice be? And can the question of permitted error be divorced from the question of the choice of scale factor?

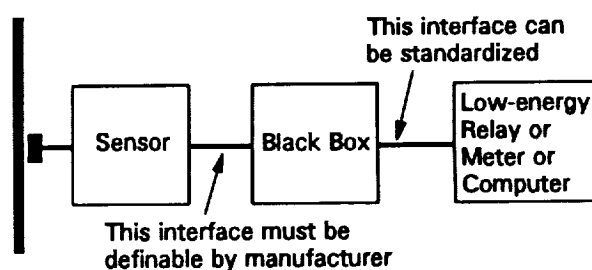
I am convinced that there are indeed two separate questions. Suppose that an interface is defined in which rated input produces 5 mA output, instead of 5 A. The sensor accuracy can be specified by a re-written version of the specification in which 5 mA replaces 5 A. The allowed error has the same meaning, in terms of percentage. At this point, however, there would seem to be no need to continue to connect the *scale factor* to the

error limits. A device can be made 1% accurate whether the output is 5 mA or 5 A, or even 5 V.¹⁸ I therefore assert that the accuracy specification can be re-written independently of the choice of the interface standard. Neither Figure 3-1 nor 3-2 depends on the fact that the secondary current is 5 A. Provided that the choice of interface does not make it difficult to meet the required accuracy, we are free to make an independent decision.

3.5 The Black Box Approach

The “Black Box” approach, shown below in Figure 3-3, is one way of addressing the question of a standard interface. A conversion device (of a nature yet to be determined) is interposed between the transducer, which we can assume for the time being to be an OCT, and the final output.

Figure 3-3
The Black Box approach
to the interface



It will be seen in this diagram that there are actually two interfaces. One is between the sensor or transducer and a device labeled black box, the other between the black box and the utility's relaying, metering or computing equipment.

The first of these interfaces, the one between the transducer itself and the black box, is entirely defined by the manufacturer of the transducer. Thus, if the transducer is an optical one, the interface is an optical interface. If the transducer is electronic, the interface could be radio or electronic or an isolation transformer. The interface between the measurement and the black box is entirely under the control of the OCT maker. Therefore, a standard at the other interface does not hinder evolution of the technology, whether optical or not. Further, there is no need for the manufacturer to reveal exactly what happens at his end of the system, so information can be kept proprietary. This interface is not the one that we want to address.

The question of interfacing the black box to the utility's equipment is the important one. The black box becomes the responsibility of the sensor manufacturer to produce. But, because it will be built to a known standard, it can be interfaced to any other of the utility's equipment built to the same standard.

3.6 Defining the Black Box

The black box has the effect of separating the measurement from the application, and the interface. This is not a question that has had to be addressed in the past, when the black

¹⁸Verifying its calibration is another matter. This is addressed separately, in Appendix A.

box consisted of only a pair of wires. It was always fair to assume that what you put in at one end came out at the other. Now, with a black box between the measurement and the application, it is important to know that what you put in is not degraded. What you put in is an optical version of the current, modulated or encoded in some suitable form. What you get out may be optical or electrical (yet to be determined) encoded in some standardized form.

At an EPRI/NBS meeting in 1987, one of the speakers observed that "the problem to be solved is not one of reproducing a 5-amp signal, it is one of getting information about a current into a computer." (Hebner, 1987.) Though it passed without comment at the time, this is an important and forward-looking statement. There may eventually be no need to specify an analog interface. There is a good possibility that much of the sensing and measurement of the future will have a digital output, so that to force the sensor back into the analog domain simply to connect it to another digital device seems to be redundant. That having been said, the simplest approach to solving the definition of the requirements for the transducer interface may nevertheless be to express them first in analog terms, such as bandwidth and dynamic range. Maybe we think in analog terms.

How to make sure the interface does not degrade the measurement? The questions start with questions like the following: What is the frequency response of a current transformer? What dynamic range is required of the measurement? From here, you have to ask, How are these parameters affected by the interface, the wiring and connections between the transformer and the relay, or by the impedance of the relay?

It seems likely that the requirements for the old high power 5-A standard or 120-V standard on instrument transformers were never examined this way. ANSI C57.13-1978 says, on the matter of frequency response, that "the frequency shall be 60 Hz." IEC 185 merely observes that the accuracy requirements are applicable at rated frequency.

3.7 Dynamic Range

Let us spend a moment looking at dynamic range. In ANSI, the accuracy of a CT is typically specified at rated primary current and at 10% rated current. This implies a dynamic range of 20 dB, very much in line with our old acoustical recordings! Voltage transformer performance is specified at 90% and 110% rated voltage, a dynamic range (if it can be called that) of less than 2 dB. Why should we be concerned with dynamic range, since the devices presently used essentially don't have one?

There are several reasons for at least considering the question. First, if we can describe a system with a truly enormous dynamic range, it may be possible to fix the effective transformer ratio for all CTs. Second, if we have to resort to multiple permissible transformer ratios, it might at least be possible to keep the number small if the system dynamic range is large. Third, if we can improve system performance without incurring a significant cost penalty, why not!

The first reason is not as outlandish as it might seem. Currents of *any* magnitude can be described by some codes; ASCII is an example. ASCII can say "1 μ A" as easily as it can say "1 MA." (It just did, in typesetting this page!) Even if we adopt a more conventional approach, the ratio of the smallest current that might interest us to the largest we would expect to encounter is probably only 60 dB, the ratio of 50 A to 50 kA. Add to

this the few additional dB required to satisfy the standard for a given measurement, and we still may have something achievable with one definition.

Our third justification for worrying about dynamic range is that the devices we routinely use are rather better than the standards would imply. At high frequencies, a typical current transformer can be thought of as a perfect transformer behind a single-pole low-pass filter, with a corner frequency in the order of 20 kHz. (This is the bandwidth of a stereo system.) And as a sensor, it is capable of good performance over a much wider dynamic range than 20 dB, often 40 or 50 dB. Add *this* to the range of currents, and you have a range that probably cannot be handled by an analog system without changing ranges.

3.8 Frequency Response

The frequency response available from optical current transducers is limited in two ways. At very high frequencies, unwanted effects like piezo-electricity cause measurement errors. It is usual to limit the frequency response in the electronics to much lower values. Nevertheless, it may be prudent, in terms of setting a forward-looking standard, to recognize that modern optics and electronics are capable of better performance than iron-based measurements. However, it should not be forgotten that dynamic range is limited, ultimately, by noise, and that most forms of noise increase with the bandwidth of the system. In other words, it is very hard to maintain a wide dynamic range and a good frequency response at the same time.

It is hard to justify defining an interface standard that maximizes performance. Thus, we should avoid specifying that "the accuracy shall be maintained up to 1 MHz" for example. The performance should not limit the application of the data to any lower performance than is presently available, but it need not go greatly higher. Perhaps all that is needed is a specification of this parameter in the device description.

The current transformer itself can reproduce signals with a bandwidth in the order of 20 kHz. To reproduce such a signal, then, the interface must not be responsible for degrading the information. To maintain an acceptable phase error at 20 kHz, the interface bandwidth has to be somewhat greater. The actual need for this bandwidth will depend on the application. There are no standards now that ensure that a given CT can be used to monitor the chopped waveform produced by a switching power supply, for example. And there is no standard for measurements below 0.05 p.u., even though some loads spend a good deal of time at this level.

Burdens are described separately for metering and for relaying. They are specified (R and L) only at power frequency. The transformers themselves are classified according to the voltage at which the output current has fallen by some amount (usually 10%) from the calculated value assuming linear transduction. Remember, the old standards were written when the transfer curve was the "lazy-S" of the hysteresis of iron. Clearly, the standard forces excellent precision on the metering CT, but recognizes the practical difficulty of designing a CT that will work in a linear fashion over an arbitrarily large range of currents. On the other hand, we know that the optical transducer will work over a moderate range of currents, and can be made stable enough for metering purposes. By considering the interface definition to be similar to a stereo system in terms of requirements, we are actually expanding on the state of the art. By going slightly beyond a mere

scale factor definition, the way is open for a range of new possibilities.

The whole process of considering all the factors and agreeing on a new standard could take years. It is important to get started now. The persistence of the existing standard for 80 years implies that it is important to do a thorough job. The notes that follow are intended to stimulate discussion, rather than to answer all the questions.

3.9 Requirements

A new interface standard should define a reasonable way of transmitting current information from a CT, or voltage information from a PT, in the usual power system environment (wide range of temperature, moderate electromagnetic interference).

The existence of a standard should enable manufacturers to make compatible equipment, and allow utilities some second-source possibilities. The standard itself should not degrade the accuracy of any measuring system, and it must not limit technology development. It should not drastically increase the price of an OCT, and it must not force manufacturers into revealing proprietary information. We have seen that these goals can be met by the Black Box approach. But what is supposed to come out of the box?

The analog approach has the advantage of simplicity. The measurement is of an analog quantity, after all, so this approach should require little in the way of special signal handling. However, we really have to go beyond the mere specification of a scale factor. There is a danger that limiting the specification to scale factor alone would lead to an interface that failed to meet the requirements discussed above. Wider consideration than the isolated choice of a scale factor is called for.

Analog signals have limited dynamic range (which leads one to suspect that more than one ratio standard may be required to cover all applications) and are susceptible to noise interference. The dynamic range limitation is present even if you use some scheme such as FM on a high frequency carrier to reduce noise susceptibility; you can help the noise immunity, but you cannot improve the dynamic range.

Digital interfacing has, typically, the advantage of immunity to noise interference, and can handle an arbitrarily large dynamic range. Further, it is likely that the signal will have to be digitized somewhere, so why not at the black box? On the other hand, digital systems are complex, and the sampled nature of the digital representation leads to the possibility of skew between data samples, which can cause problems in some calculations.

A summary of the advantages and disadvantages of various approaches is given in Table 3-2. A detailed examination follows a brief aside to consider one more factor.

3.10 Additional Considerations

So far, we have considered the question of the interface as if the situation was always one where *one* CT was connected to *one* relay, or meter. This is the situation implied by Dr. Hebner's comments in Gaithersburg. More generally, the CT will be part of a system, involving other CTs, and multiple relays.

What this means is that one CT may be used to furnish information to a local relay in the station yard, and the same information to another system in a control house, perhaps

several hundred meters away. With 5-A CTs, it may also be that the output of several CTs may be combined to drive one relay. And if the output of more than one CT is used as input to a protection system, some degree of “matching” may be required.¹⁹ If the CTs are digital (or if the interface is digital), there may be a need to synchronize the sampling periods. Our proposed interface standard must not preclude any of these possibilities.

Table 3-2 Comparison of Interface Options

Option		Advantages	Disadvantages
Digital	Delta	Easy to transmit optically. Is self-synchronizing. Is fairly simple to implement	Care is needed to achieve high enough stability for metering. Small dynamic range may result in need for multiple scale factors
	Parallel	Allows easy definition of scale factor. Unlimited dynamic range: one scale factor could serve all needs	Difficult to implement optically, but susceptible to noise if not optical
	Serial	Easy to transmit optically	Some data skew is inevitable. May require parallel channels in high-performance versions
	FM	Easy to transmit optically. Simple to recover analog signal	Dynamic range may not be large enough to allow the use of only one scale factor for all CTs
Analog	Parallel	Easy to implement. Permits loads in parallel	Limited dynamic range because very susceptible to noise pickup in wiring
	Serial	Easy to implement. Permits loads in series	Limited dynamic range because moderately susceptible to noise pickup in wiring

It is most likely that the output of one CT will have to furnish information to a number of locations. This is certainly attractive from the cost point of view. With a 5-A CT, it is a simple matter to connect the various relays in series, provided the total burden limits are not exceeded. We might allow for the same possibility with our low-energy

¹⁹It is quite possible that the lack of any specifications for dynamic range and frequency response has resulted in protection system failures when non-matched CTs are used together.

interface. Another frequent variation on the one-CT, one-relay theme occurs when three relays are connected in parallel. In many applications, the sum of the three phase currents is zero except during fault conditions. These simple features do not work in the optical domain. How will this influence our interface?

If we specify an interface that is implemented optically, it is certain that the signals in the fibers cannot simply be added to generate a "sum." Except for the carefully controlled outputs of some optical sensors, the optical power level is not particularly relevant. And anyway, when you add optical signals, all you get is an interference pattern. So much for adding sensor outputs. As far as driving multiple burdens is concerned, whatever the signal representation is, it would be unwise (and expensive) to tap a fiber to serve multiple loads. Therefore, any optical interface we choose carries the burden of requiring conversion to the electrical world before even the simplest processing can be done on it. But provided all the signals are according to the same standard, they can easily be added (or subtracted, or multiplied...) or repeated as electrical signals, either digital or analog.

Let us now turn to the question of digital or analog.

3.11 Digital Interface

If the interface is digital, the number of options is great. We can have serial or parallel data, and we can represent the signal in a variety of ways. One possibility that I think should be considered is the representation of the measured quantity in real units (like Amps) rather than in per unit.

Consider: a sixteen-bit integer representation can handle 65,535 A ($2^{16}-1$) with a resolution of one amp. Note that a 16-bit specification for the interface does not require the use of a 16-bit A/D convertor. Suppose you wished to use an 8-bit A/D with an instrument rated at 2000 A and leave a factor of eight for "headroom." If the system were parallel, the 8-bit A/D convertor would be connected to bits 6 through 13 of the interface (see Figure 3-4). Bits 0 through 5 would be zero because they were not used. Bits 11 through 13 would be available, but would normally be zero because they represent overcurrent. Bits 14 and 15 are not used. The system resolution at any single sample would be 64 A.

This kind of approach would lend itself very well to a parallel digital interface. This interface has the advantage that, for any given bit rate, the delay between the sampling instant and the reception of the data word is small. However, a multi-bit parallel digital interface is cumbersome, and is difficult to transmit optically. A serial digital interface is ideal for optical transmission, but necessarily involves some complexity to encode and decode.

3.11.1 Sampling

The question of how many bits per sample, and how many samples per cycle—beware of confusing samples per cycle and samples per second—is highly dependent on the application. EPRI report EL-1601 (Schweitzer *et al.*, 1980) describes a revenue metering

system that takes only 4 (12-bit) samples per cycle to exceed the precision required for billing. The ASEA digital OCT developed by Adolfsson's group takes rather more: about 16 samples per cycle (Adolfsson, Einvall, Lindergh, Samuelsson, Ahlgren and Edlund, 1989). Although it is sometimes stated that a faster sampling rate is needed, or that a 16-bit converter must be used, it is by no means clear that these extremes are justified.²⁰

Figure 3-4
Possible digital representation
of current

Bit Number	Current (A)		Bit Number	Current (A)
0	1		0	64
1	2		1	128
2	4		2	256
3	8		3	512
4	16		4	1024
5	32		5	2048
6	64		6	4096
7	128		7	8192
8	256			
9	512			
10	1024			
11	2048			
12	4096			
13	8192			
14	16,384			
15	32,768			

A/D Converter

Interface Specification

The resolution problem with an A/D converter that has few bits is not as serious as it might seem. In the example I gave earlier, a resolution of 64 A resulted from the use of an 8-bit converter. The obvious, brute-force method of improving the resolution would be to add bits. A/D converters with 12 bits are fairly commonplace. Converters with 16 bits can be made, though they are expensive. The use of either would improve the system dynamic range. There is an alternative, however. The quantization error (64 A in our example) has a statistical distribution (the *expected* error is zero) that permits a simple averaging process to reduce its effect.

This is a simple form of a digital signal processing technique known as "noise shaping," that is now finding application in compact disc players (for D/A conversion). The approach is based on delta modulation, which calculates a signal according to its change from the previous sample. (Appendix B describes delta modulation. Appendix C describes noise shaping.)

Consider the *information* contained in a sampled signal. If a signal is sampled in two different ways, one with twice as many samples per second as the other, but with one less bit per sample, there is no difference in the amount of information the two representations contain. I said above that a CT had a dynamic range of 40 or 50 dB at best. It happens that an A/D converter of 8-bit resolution will give a dynamic range of 48 dB. So if a converter of this quality is used, sampling at the Nyquist rate, the signal will contain all the

²⁰The device described by Adolfsson *et al.* uses a 10-bit A/D converter in combination with a prescaler to achieve a dynamic range equivalent to a 17- or 18-bit converter. This performance is justified by the desire to meet the requirements of both metering and protection in a single system.

information available at the output of a conventional CT.

If a CT really has the frequency response of a lag filter with a corner frequency of 20 kHz, it could be completely represented by sampling at 44 kHz, with 8 bits per sample. (44 kHz is a standard rate that happens to be just slightly above the Nyquist rate for the device in question.) As a serial data stream, a channel with at least 352 kb/s capacity would be required for communication. (In practice, synchronizing and error correction would increase the requirement.) For every bit added to the converter, the data rate must be increased. At 12 bits, the rate is therefore 528 kHz, and at 16 bits the channel speed requirement is over 700 kHz.

44 kHz and 16 bits are the numbers for one of the two channels on a digital compact disc. The specifications for CD make it very clear that low distortion, wideband sound reproduction is achieved. The linearity and frequency response far exceed any needs presently met by CTs. It is therefore urged that any statement that might be made to the effect that more than 8 bit samples, or more than 44 ksamples/s *are needed* for a CT interface be viewed with skepticism. There might be cases, in research for example, of applications that require measurements of greater bandwidth or dynamic range, or sometimes both. But such performance far exceeds the performance of CTs developed from designs made 80 years ago, when current transformers were hand wound, with wire insulated with silk.

3.11.2 Problems with Digital Interface

While parallel digital is conceptually simple, in practice it can be cumbersome in terms of cabling. Further, the multiplicity of wires leads to some increased susceptibility to interference. The fact that serial data can so readily be transmitted optically argues in favor of that approach. An optical digital signal can be widely distributed inside a switchyard, without concern for interference.

There are two problems with digital serial. One is that there is necessarily a delay between the sampling instant and the reconstruction at the receiver of the complete data word. The other is that serial data must necessarily be encoded and decoded: start and stop bits must be added, as well as error checking bits (such as parity) and error correction bits. The Philips/Sony serial digital standard used in CD players might be used, because that would be suitable in terms of frequency response and dynamic range. However, the complexity and inherent time-lag of that approach may rule it out.

Bear in mind that the applications for OCTs, and therefore for this interface specification, include high-speed large dynamic range measurements for relaying for high-voltage systems, high-precision measurements for billing in high-voltage systems, and low-cost low-quality measurements for the distribution system. It seems as if we have a set of incompatible requirements. If the interface has to be made as good as the most stringent application demands, this may be an expensive proposition, which would be unfair to the manufacturers of low-cost hardware for the distribution system.

There may be a way round this problem. I will address that next.

3.11.3 A Modified Black Box

The expense of serial digital lies in the need to encode additional information such as framing data and error codes with the signal at the transmitter, and decode it in the receiver. If a low-budget system is designed in which the information is not actually required, as in a distribution system application, this is an unwarranted expense.

The solution is to adopt the philosophy of color TV, or FM stereo: start with a minimum system, and add features that do not necessarily have to be used. Thus, if you want to save a little money, you can buy a monochrome TV, or a mono-only receiver. The authorities that regulate such things (the FCC in the U.S.) have seen to it that the more complex signal being transmitted is compatible with the simpler receiver.


In our situation, this approach would mean that we start with a simple digital serial interface, and add features such as error detection/correction coding, or synchronization inputs, in a parallel channel.

Strictly speaking, from a cost point of view, it is time-division multiplexing that we should avoid. This way, low-end systems are not faced with the cost of stripping unused information out of the data stream. Color TV and FM stereo are frequency-division multiplexed, essentially in parallel in the frequency domain. In an optical system, we have the additional (but not inexpensive) alternative of wavelength-division multiplexing.

However, we are not forced to multiplex our signals. It would likely be more cost effective to design a hybrid interface of several serial connections in parallel.

Although I have expanded the representation of the OCT-to-black-box interface in the diagram (Figure 3-5) to show two-way information, the details are not our concern. The black-box-to-user interface, here shown as four lines, is the one to specify.

Figure 3-5
Modified Black Box interface



As well as the output signals, one or more input signals may need to be defined. For example, the input line might be needed in order to synchronize the sampling instant at a number of devices. We might end up with two or perhaps as many as four separate serial connections: data out, error coding out, timing information out; synchronizing signal in. In a high-performance system, the full set of information might be used. In a low-budget system, perhaps only the data output has to be implemented.

3.12 Analog Interface

If the interface is analog, should the measured quantity be represented by a voltage or a current? Either is easy, and both are appealingly simple at first sight.

Suppose, for an example, that we decide to represent rated primary current by a signal of 100 mV. ANSI C57.13 talks about the accuracy of a relaying CT at an overload

of 20 times. (A ratio error of 10% is permitted.) This is a 2-volt signal, not unreasonable for solid-state electronics. The same standard discusses the required accuracy of a metering CT at 10% load. In our example, this would be 10 mV. There's a problem here. It may not be reasonable to expect a 10-mV signal to be transmitted faithfully in the electromagnetic environment of a high-voltage switchyard or a power station. And for metering purposes, the volt-drop in interconnecting wires cannot always be neglected.

The situation is a little better if we use, say, a 10-mA signal to represent rated primary current. At the high end our electronics is called upon to generate a fifth of an amp, which is acceptable, but at the low end we are dealing with a 1-mA signal. This is too small to be taken seriously if it has to be transmitted any distance in station environment.

3.12.1 Problems with Analog

While the complexity of the digital world argues in favor of an analog specification for the interface—and at least one manufacturer has opted for this solution—there may be problems. Power stations and substations are often noisy places, especially at critical times. Wiring from an OCT black box in a switchyard to a relay in a control house would be susceptible to interference, and the relay would be prone to false operation. The only way an analog interface can solve the noise/dynamic range problem is for the interface wiring to be very short. Ideally, one would locate the black box inside the relay housing.

It is important to recognize that this is a non-solution. **If the black box has to be inside the relay box, why have a black box at all?** You might just as well make a relay/OCT combination, with dedicated hardware. The very purpose of the communication standard we are trying to develop is to avoid the need for measuring and relaying systems to be bought as one piece. By defining the properties of information exchange at the physical boundary of the hardware, we allow for independent development of the hardware on both sides of the interface. It seems that the noise problem of the direct analog interface just about precludes that.

Delta modulation represents a class of digital communication means that might be termed "indirect." (I class it as digital because the amplitude is fixed and the timing of the output pulses is quantized by a clock. I classify FM as analog because the instantaneous frequency is not quantized, but can change infinitesimally.) Indirect analog communications, such as frequency modulation, are capable of handling signals over a moderately large dynamic range, and can be designed to have a frequency response that will cover any reasonable power system requirement. The resulting system is relatively simple, even though careful design may be needed to achieve the required accuracy over the temperature range usually demanded of power system components.

I regard this option as an excellent one for our application. The system is *fairly* simple, and lends itself to serial transmission, ie, it adapts readily to an optical system.

3.12.2 Frequency Modulation

The use of an FM carrier is an approach that might meet all our needs. With an FM system, the carrier frequency can be chosen to simplify filtering, and the signal can be

amplitude limited before demodulation, to help preserve the noise immunity. With a reasonable depth of modulation, the linearity of such a system can be excellent. The method also lends itself readily to optical implementation, and in this form the noise immunity is excellent.

The dynamic range of an FM system can exceed, say, 65 dB, if suitable care is taken in the design. FM is not particularly complex, and needs only a modulator, a demodulator and some kind of filtering, so the advantage of extreme simplicity is not greatly reduced.

The filtering requirements arise from the need to reduce the carrier energy remaining in the analog signal after demodulation. The spectrum of the demodulated signal typically contains strong lines at the carrier frequency and its harmonics. The filtering requirements can be simply estimated. A 2-pole filter rolls off at 40 dB/decade, so that a frequency ratio of 1000 allows for a theoretical amplitude reduction of 120 dB. In practice feedthrough, noise and other considerations limit the feasible reduction to 60 or 70 dB.

It would be only moderately difficult to make such a scheme work well enough, over the usual temperature range, for a metering system. Quite ordinary components can be used, and the stability can be made arbitrarily good. This is the approach we have used at JPL in the telemetry for our field meters, achieving better than 1% accuracy (after temperature compensation) over a range of 80°C.

3.13 A Perspective

We have talked about systems that have a dynamic range of 100 dB, five orders of magnitude. We have mentioned sampling rates in the megahertz region. It is important not to lose sight of the fact that today's systems perform well enough with limited specifications. The information presently available from a current transformer has a frequency response that may contain no useful information above 1 or 2 kHz, has a dynamic range of only a few dB, and is transmitted successfully over circuits that would be scarcely adequate for a telephone.

We would not want deliberately to degrade the performance of our proposed low-energy system to this level, but we should be prepared to ask *why* if we start to work on an expensive interface simply because it has greatly improved specifications. If the job does not call for anything remarkable and new, improved performance is not justified.

When I started to prepare the material that went into this section of the report, I had no favorite solution. In the process of putting my ideas in order, I have become persuaded that there *is* one approach that might well be good enough to meet all our reasonable requirements, and is almost as simple to implement as the direct analog interface it replaces. Frequency modulation with a carrier frequency in the order of, say, 10 MHz, would be quite straightforward to implement in an electrical or optical version. Inexpensive optical components (such as LEDs rather than lasers) could be used in an optical version. Since the highest modulation frequency of interest is in the order of 1000 times lower in frequency, filtering residual carrier energy from the demodulated signal should be simple.

It is uncertain whether the fact that the demodulated output signal is analog should be regarded as an advantage or a disadvantage. Within the confines of the interface device, where it can be shielded from external interference, it is probably not a disadvantage. The fact that analog signals lend themselves readily to interfacing directly with existing metering

and protection devices that use low level signals must be considered an advantage. Analog signals can also be adapted to the simple kinds of arithmetic required of CTs: typically no more complex than summation.

In an attempt to understand and formalize my preference for FM, I have developed a simple spreadsheet that enables the quantification of the subjective judgments of what is important in the interface specification and how well the requirements are met by the various options. This is described next.

3.14 Figures of Merit

One way to compare the alternative implementations of an instrument interface is to derive for each a Figure of Merit. The approach shown below simplifies the procedure by separating the requirements from the abilities of the implementations to meet them, and automating the calculation.

Suppose there are n separate factors that qualify as requirements. Examples could be dynamic range, noise immunity, and ease of interconnection. These could be ranked in importance by assigning each a score or weighting, say between 0 and 10. Together the rankings could be written as a matrix R of dimension $1 \times n$.

Now suppose that the performance of a particular interface is judged as to how well it meets these various requirements. A parallel digital interface, for example, might score high in terms of dynamic range and noise immunity, but low in terms of ease of interconnection. The opposite would be true for a serial analog representation. The overall value of the two approaches can be compared by multiplying the score each achieves in meeting a given requirement by the weighting of that requirement, and forming the sum of the products.

Each implementation approach can be considered to have a performance matrix P , of dimension $n \times 1$. By expressing the requirements and performance in this way, it is clear that a product RP can be formed that results in the required scalar quantity or Figure of Merit, $FOM = \sum R_j P_j$.

Although spreadsheet programs are not particularly adept at doing matrix algebra, the required expressions can be programmed into the appropriate cells to form a handy, interactive calculator for Figures of Merit. The results of such a calculation are shown in Table 3-3, which of course reflects the biases of the author.

There may be one alternative worthy of serious consideration. Recently, the world of hi-fi has put a new optical fiber interconnection on the market. I have (so far, and despite considerable effort) been unable to ascertain its specifications, but it is inconceivable that in most respects it would not easily do our job, at minimum cost. Stereo is a far more demanding application than ours: and the demands have evidently been met at low cost. (How much would you spend for a stereo?) Though I now favor FM, I am trying to find out more about this interface.

Table 3-3 Printout of spreadsheet for Figures of Merit

REQUIREMENTS WEIGHTINGS						
2	Dynamic Range					
3	Frequency Response					
7	Cost					
9	EMC in yard					
5	Multiple loads					

	IMPLEMENTATION SCORES					
	serial	ANALOG parallel	FM	serial	DIGITAL parallel	delta
Dynamic Range	5	5	5	9	9	6
Frequency Response	5	5	5	4	6	4
Cost	8	8	7	4	3	5
EMC in yard	2	1	10	10	10	9
Multiple loads	7	7	5	5	5	5
FIGURE OF MERIT	134	125	189	173	172	165

3.15 A Process

The question of the application cannot be completely separated from the question of the interface, even though this is suggested by the Black Box approach. Nor is it likely that all the issues have been addressed satisfactorily in this report. One way the matter might proceed is outlined in Figure 3-6.

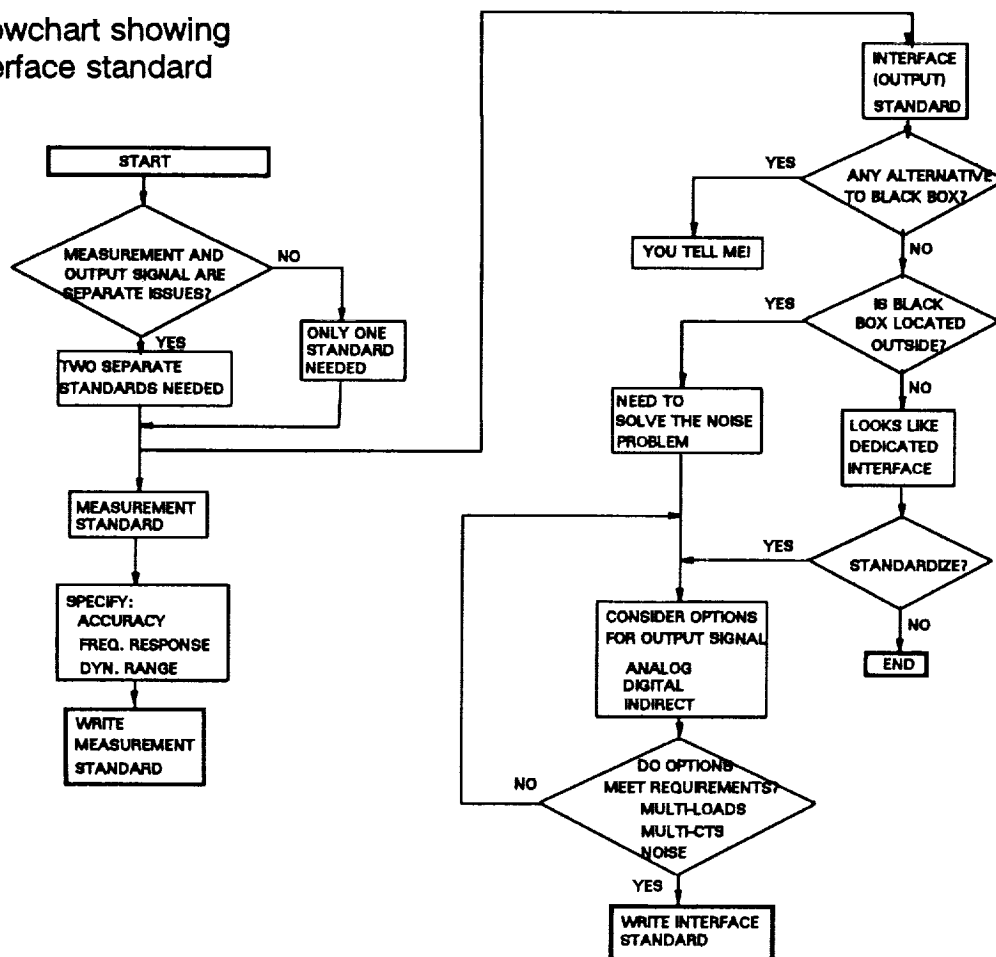
If agreement can be had on the need to create a new standard, and on the usefulness of the black box approach, then the selection of the interface (shown as the last two blocks in the Figure) might proceed as follows:

- **The requirements are studied.** This work takes into account the various applications (such as billing, protection, operation) and the features of the existing high-power standards (such as their accuracy and noise immunity). From this is synthesized a set of interface requirements in terms of accuracy, dynamic range, bandwidth, noise immunity, maximum time-lag etc.
- **This set of interface requirements is publicly criticized.** It is issued as a paper by IEEE PES. Comment is solicited in committee meetings. Other organizations, such as the IEEE Standards Coordinating Committee on Photonics (SCC-26), are involved.
- **Ways to implement the requirements are studied.** (Likely candidates should be included in the IEEE paper, so that comment can be sought.) There are many standards already in existence that might be suitable, but a new one may be required. Ultimately, suitable implementation definitions are issued as IEEE guidelines, then recommended practices, then Standards.

The first step, the requirements definition, is crucial. No such study seems to have been undertaken, as far as I know.

It is, I believe, important that any replacement for the old standards should be very carefully considered. The manufacturers of relaying equipment as well as instruments should be involved. Without doubt, the work would benefit from utility involvement. It has to be very clear to all concerned that devices built to the new standard perform as well as devices that were built to the old standard, even if the performance requirements were never previously identified.

Figure 3-6 Flowchart showing route to an interface standard



3.16 Concluding Remarks

This section of the report has examined some of the ways a low-energy measurement transducer could be interfaced to relaying or metering equipment. The options are many. At the broadest level we must choose between analog and digital; serial and parallel.

If we select digital we must answer further questions regarding the number of bits per sample, and the number of samples per cycle. If the choice is analog, can the equivalent questions be ignored: the dynamic range will be affected by the choice of nominal scale factor, and the frequency response could be limited by the electronics. No clear answer seems to have emerged yet.

If we are successful in defining a new interface, low-energy instruments may someday be responsible for the demise of the 5-A standard.


Not before time.

APPENDIX A

CALIBRATION OF CURRENT TRANSFORMERS

THE COMMENT is often heard that optical current transducers will be hard to calibrate. Why should this be? What is so different about an optical CT compared to a conventional CT that will make it difficult to calibrate? After all, a simple set-up like the one shown in Figure A-1 should surely do the job.

Figure A-1
Straightforward method for
current transducer calibration



The primary current is measured, presumably by means of a standard CT as shown, or perhaps using a shunt resistor. The output of the test CT is measured with a good voltmeter or ammeter. The method is simple. The problem is that all the uncertainties of the individual measurements add to the calibration uncertainty. Thus, even if the primary current is measured through a standard CT, the uncertainty of the standard CT is equivalent to uncertainty in the value of the primary current. The situation is similar as far as the test CT output is concerned. The OCT calibration uncertainty therefore includes the uncertainty in the standard CT, the uncertainty in *measuring* the output of the CT, and the uncertainty in measuring the value of the output of the device under test. Nevertheless, Papp and Harms (1980) used an approach similar to this to investigate the errors in a magneto-optic current transducer, and made measurements with resolution an order of magnitude better than the error limits of IEC 185. The overall calibration is effectively done with respect to the output of the standard CT, however, and errors in this device will lead to undetected errors in the calibration.

Cease and Johnston (1990) report long-term experience with a magneto-optic current measurement in an outdoor substation environment. The data from the optical measurement are compared with data obtained at the same time from a conventional current transformer. There is evidence of very slow drift, of small magnitude, between the two measurements, and this has been attributed to the effect of temperature on a seasonal time scale. The problem is that there is no way of distinguishing drift in the optical transducer from drift in the current transformer.

Figure A-1 is not a diagram of how calibrations are usually done. In general, methods for calibrating measurements or units can be classified into two groups: *direct* and *comparator*. Within these groups, the calibration can be further subdivided into *null* methods and *deflection* methods. Direct methods do not require a standard of the unit being calibrated. An example of this approach is the Lorenz method for determining resistance, which fixes the resistance in terms of parameters such as length and rotational speed. Similarly, current can be defined in terms of the force between two current-carrying conductors: the Rayleigh force balance embodies this principle. A comparator method requires another standard: it used to be that length was determined with respect to a meter bar, for instance. Mass is still determined by comparison with a standard mass, kept in Sèvres, in France.

Deflection methods rely on the indication on some instrument, whereas null methods are designed such that the only point on the indicating instrument that has to be known accurately is the zero. This has seemed, over the years, to have enormous attraction, and has given rise to many elegant techniques. The Lorenz method of current measurement is a (direct) null method. The *bridge* methods of measurement are all (comparator) null methods.

A comparator method for calibrating a current transformer requires, of course, a standard current transformer (or other similar network) with at least a similar ratio to the one being calibrated. ANSI recommends the use of a current comparator method as being capable of the highest measurement accuracy, with uncertainties in the order of a few parts per million. Consider the arrangement shown in Figure A-2.

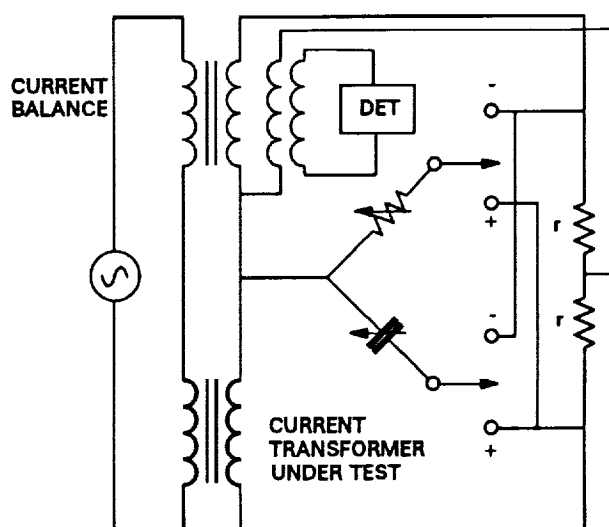


Figure A-2
Current transformer accuracy test
with current comparator

The method is a comparator/null method. The current comparator is basically a current transformer operated with no net flux in its core, so that many of the defects of the magnetic circuit are eliminated. It consists of a toroidal core with a number of windings. Nearest the core is the detector winding, whose purpose is only to detect the condition of zero flux in the core. This is covered by a compensation winding. There follows a magnetic shield, essentially a concentric magnetic core for the current transformer with its primary and secondary windings on the outside.

If the circuit is balanced, there is no flux in the current comparator. Neglecting the compensation winding, this means that the primary ampere-turns and the secondary ampere-turns cancel. There is no volt-drop across the current comparator, so the burden seen by the test CT is not affected by it. If the test CT is not perfect, a small error current will flow in the secondary circuit, and result in a flux in the comparator core. This is canceled by the compensation winding, which carries a current derived from an RC network. (The switches shown in Figure A-2 allow positive or negative errors to be compensated.) At balance, the error in the test transformer is given by

$$\varepsilon = \pm \left[\frac{r}{R} + jr\omega C \right] \quad (\text{A-1})$$

where the real term gives the ratio error and the imaginary term the phase angle.

Several important differences from the method of Figure A-1 emerge.

- The method is a null method. No quantity except zero has to be measured precisely in order for the device under test to be calibrated. Indicating instruments are usually accurate at zero.
- Not even the primary current has to be measured accurately. It might be reassuring to know that the calibration was done at some particular current, but knowledge of the current is not essential to the calibration.
- Uncertainties in the measurement of the resistors shown in Figure A-2 contribute only to uncertainties in the measurement of the error. In other words, these become second-order effects. If the two resistors are 1% components, and the error can be given as, say, 0.1%, then the ratio is known with an uncertainty 50 times less than 0.1%, or 20 ppm.

There are a number of variations on this theme. In one, an auxiliary transformer is used with the current comparator. In another, a standard current transformer is used along with the current comparator. There is also a direct-null method, a simplified version of which is shown in Figure A-3, in which the ratio is expressed in terms of the ratio of two resistors, and the phase angle in terms of a resistance and a mutual inductance. (Mutual inductance standards can be designed such that the value can be calculated from the dimensions alone.)

The detector is nulled by adjusting R_2 and the mutual inductance M . The ratio is then given by R_2/R_1 , and the phase error (in radians) is very nearly $\omega M/R_2$. Since it is the ratio, rather than the error, that is found, the method is less forgiving than the current balance of Figure A-2.


Each of the various methods mentioned above has its own advantages and disadvantages. Some methods are viewed as inherently capable of greater accuracy than others. Some techniques impose limits on the test CT burden. All require the generation of a large

primary current if the calibration is to be done under realistic conditions. Recognizing that the main source of error is the magnetizing current, the ANSI and IEC specifications permit a different kind of test, in which the secondary of the CT is energized, and the primary is open circuit. The resulting current should not exceed the limits for a parameter called the “composite error” defined by the expression

$$e_c = \frac{100}{I_p} \sqrt{\frac{1}{T} \int_0^T (K_n i_s - i_p)^2 dt} \quad (\text{A-2})$$

where the variables have the meanings given in Section 3. Unfortunately, it is not permitted to use this test as an accuracy test for metering applications.

Figure A-3
Current transformer accuracy test
with direct-null network



A.1 Applicability to Low-Energy Sensors

Can the calibration methods used for CTs be applied to low-energy sensors? An optical or electronic sensor might indeed have a current output. Considered as a CT, such a device would have a distinctly unusual ratio. A normal CT might have a ratio of 5000:5 for example, the output being 5 A for 5 kA input. An optical CT might produce 5 mA under these conditions, an effective turns ratio of 10^6 . While the number of turns in a current comparator could be counted with sufficient precision, and a million-turn current balance probably could be made, it is unlikely that a magnetic circuit with so different a pair of windings in the primary and secondary would not have serious defects. Cascading two transformers would not help: the current balance is more than just a simple transformer.


Eddy So of Canada's National Research Council (the equivalent of the U.S. National Institute for Standards and Technology) has developed current comparator techniques for the calibration of low-energy instruments, and has written extensively on the topic. In a paper presented at an IEEE Power Engineering Society Meeting in India in 1990, Dr. So reviewed the development of calibration instrumentation at NRC.

First, a circuit called a transimpedance²¹ is described. By using a current comparator approach, a circuit is developed in which the small defects (typically due to temperature dependence) in the electronics can be compensated by a feedback arrangement. The result is a circuit that produces a voltage output that is proportional to the current input, and that is very stable. The circuit is not simple, consisting of a current transformer (to reduce the current in the electronics), a current balance, and a step-down transformer (to further reduce the current in the electronics). A feedback system is used to null in the circuit output the error detected in the current balance. Defects in many of the components are reduced by the gain of the feedback circuit, so that the main sources of uncertainty due to components are the input current transformer and a standard resistor.

Complex and advanced as it is, the details of the transimpedance circuit need not concern us here because for calibration purposes it can be regarded simply as a surrogate optical current transducer. The problem we are trying to solve (and indeed the problem So was addressing) is not how to build it, but how to calibrate it!

The calibration circuit used by So is based on an additional current transformer and an additional current balance. The current balance is used to compare the primary current with a current derived from the output voltage of the transimpedance circuit by means of a reference resistor. (A capacitor is used to correct phase defects, but these can be neglected for the time being.) A simplified version of the calibration arrangement is shown in Figure A-4.

Figure A-4
Calibration circuit
using current balance



At balance, the ampere-turn contributions from the CT and from the transimpedance circuit being calibrated cancel one another. The detector null depends therefore on the ratio of the CT, the value of the reference resistor, and the turns ratio of the windings in the current balance itself.


²¹The IEEE Dictionary (American National Standards Institute, 1977) allows the use of the word *impedance* to mean either the ratio of the voltage to the current in a circuit, or the physical device that causes the effect. It does not mention the same dual meaning in association with the word transimpedance, which is defined as the ratio of the voltage at one pair of terminals of a network to the current across a different pair. I shall therefore use the term *transimpedance circuit* to denote the physical device.

Note that these are all first-order effects. The circuit compares the ratio of the device being calibrated with the ratio of a standard CT. The ratio of the CT, and the value of the reference resistor are external to this process, and their uncertainties add directly to uncertainty in the calibration of the ratio of the device being tested.²²

It is my opinion that, in spite of the fact that the calibration approach retains the use of a current balance, it retains only a few of the advantages. It is true that the value of the primary current is not measured, but there are still two components, the CT and the reference resistor, contributing directly to the calibration uncertainty. Only the fact that the calibration is a null method makes it preferable to the brute-force approach of Figure A-1.

It seems that a simpler, non-magnetic method could be devised that is at least this good. A bridge approach might offer a solution. Assuming that the optical current transducer produced a voltage output (as does the transimpedance circuit), one possible implementation is shown in Figure A-5.

Figure A-5
Bridge-resistor method for calibrating
optical current transducer



Although the method indicated in Figure A-5 is intended to be a null method (the output voltage V is not used in the calculation), the detector is more than a simple null detector. Assuming that the resistor R is non-inductive, the circuit will not balance unless there is no phase error in the OCT. Unless the resistor is “padded” with inductive or

²²It is not difficult to show that if a quantity y is the product of two parameters, u and v , each of which can be measured with uncertainty Δu and Δv respectively, the maximum uncertainty in y is

$$\frac{\Delta y}{y} = \frac{\Delta u}{u} + \frac{\Delta v}{v}$$

A similar expression holds for quotients, and for extended products. However, if the various uncertainties are uncorrelated, a statistically reasonable estimate of the uncertainty is

$$\frac{\Delta y}{y} = \sqrt{\left(\frac{\Delta u}{u}\right)^2 + \left(\frac{\Delta v}{v}\right)^2}$$


Because of the squaring of the individual contributions to the uncertainties, the largest will tend to dominate the total. For example, with four terms contributing to the total, with magnitudes 0.3%, 0.1%, 0.1% and 0.1%, the total uncertainty is less than 0.35%. A fuller discussion of this topic is presented by Holman (1978).

capacitive trimmers, the detector has to measure any resultant phase error. And since no adjustment has been shown on the resistor, the detector will have to measure the ratio error too. The method is thus not quite a null method. However, the deflection of the detector involves only a second order uncertainty as far as the OCT is concerned.

The uncertainty in the OCT ratio, rather than in the error, depends on the uncertainty in R . This will be the dominant uncertainty. Knowing the value precisely sounds simple, until it is realized that R is a very unusual device. If the OCT produces a 5-V output, the volt-drop across R must be 5 volts. If the primary current is 5000 A, R must have a value of 1 m Ω . However, it is dissipating 25 kW. This means that a precision non-inductive resistor must be used that will be stable (or whose temperature effects are well known) at relatively high values of power. In spite of the difficulties, this may be a reasonable approach because it may be simpler to calibrate a resistor to the required precision than a reference CT.

If the resistor is replaced by the combination of a standard CT and its burden, the method is analogous to So's calibration approach shown in Figure A-4, except that the current balance is replaced by an electronic detector. This is shown in Figure A-6.

Figure A-6
Bridge-CT method for calibrating
optical current transducer



If a current transformer can be calibrated with sufficient accuracy (by conventional means) with a burden connected, instead of a short circuit, then the method indicated in Figure A-6 should be at least as good as the current balance approach shown in Figure A-4. If it happened that the OCT produces a current output instead of voltage, a load resistor (reference resistor) may be used to produce a voltage output. Of course, uncertainties in the value of the reference resistor add directly to uncertainties in the calibration, but this would not render the approach overwhelmingly unattractive. Figure A-4 shows that the use of such resistors in this application is acceptable. (If the electronic output of the optical sensor produced instead a digital signal, some new way of calibrating the CT must be found.)

A.2 Digital Techniques

Digital techniques may offer further help. If the detector in Figure A-6 is replaced with an A/D converter, sampling the difference between the voltage across the resistor and the output of the OCT, it would be a simple matter to ascertain what the error in the OCT was. And since it would be the error that was measured, errors in the digital measurement would be second order effects.

However, A/D converters can be made with very good performance, particularly in terms of linearity. For example, an error in the least significant bit of an 18-bit converter represents less than 4 ppm. This justifies their use in direct measurements of the outputs (rather than the difference between the outputs) of the circuits of Figures A-4 and A-6. The output voltage V in Figure A-4, and the voltage across the resistor R could be sampled, using a multiplexed converter. A high-accuracy voltage source could be sampled, too, removing any need for extraordinary precision (as opposed to linearity) in the A/D converter. Although this is a deflection method, it is possible that these measurements could be made with sufficient precision that the dominant uncertainty in the calibration was the value of R . The advantage over the bridge method is that the value of R could be reduced significantly, thereby reducing the power dissipation, and making the resistor more practicable.

The entire calibration could then be controlled by computer, with data analysis yielding the effective ratio and phase errors. It is possible that the OCT could be further characterized, without great additional effort. For example, an impulse or switching test might be used to deduce the frequency response of the OCT, and its linearity might be deduced from measurements made at more values of current than are required by the present test protocols. If these measurements were all done in a matter of a cycle or two, the power dissipation in the standard resistor could be kept very low.

A.3 Concluding Remarks

The calibration of optical transducers, as opposed to conventional current transformers, implies that the favored current balance method might as well be abandoned. In calibrating a conventional current transformer, this method has the advantage that the ratio error is determined in a device operated with zero net flux. Only with an additional CT (not operated at zero net flux) can this feature be retained in the calibration of an optical current transducer. The uncertainty in the CT ratio and in the value of a resistor then add to the calibration uncertainty.


Making a precise, stable optical CT has been a challenge taken up in a number of manufacturers' laboratories. Calibrating it is the next challenge, presumably to be addressed in the national standards laboratories.

APPENDIX B

DELTA MODULATION

SOMETIMES known as single-bit A/D conversion, the delta modulation technique was originally developed by the telephone companies to generate a low-bandwidth serial data stream from analog data. If FM transmission can be termed indirect analog, perhaps delta modulation can be called indirect digital. It has the advantages that it is simple to implement and that signal regeneration is self-starting without the addition of framing information to the data.

In its simplest form, the method works as follows. At the modulator, the input signal is compared with the output of an integrator. If the signal is higher, the input to the integrator is such that its output tends to increase. If the signal is lower, the integrator is made to decrease. This comparison is performed at a constant rate (by means of a clock) so that the maximum data rate is simply the clock frequency. This rate should be at least twice the highest frequency of interest, but is a much lower frequency than required for A/D conversion. Figure B.1 shows the approach.



(a) Implementation

(b) Waveforms

Figure B.1 Simple delta modulation

To recover the signal in the analog domain, the data stream is applied to an integrator. This integrator tracks the one in the modulator, and the reconstructed analog output is comprised of a series of constant slope ramps in each direction.

The chief drawback of the method is its limited dynamic range (which is a function of the clock frequency). Input signals must be frequency and amplitude limited. Changes to improve the situation both reduce the simplicity of the method and increase the bit rate of the data stream.

APPENDIX C

NOISE SHAPING


NOISE SHAPING operates in the digital domain. By increasing the sampling rate, and using a kind of negative feedback, the number of bits required for a given signal-to-noise ratio at the output can be reduced. Consider, for example, the audio world. A compact disc recording has a dynamic range of 96 dB. To take advantage of this, the player must have a signal-to-noise ratio of 96 dB. This requires at least a 16-bit converter: $20 \log_{10}(2^{16}) = 96$. This is extremely difficult to do. There are two problems. Imagine that the signal at some instant is represented by a number consisting of a zero and 15 ones. A tiny increase in the signal would change the number to a one followed by 15 zeroes. For the reverse process to happen properly, ie for the recovered analog signal to be accurate, the switch from 011111... to 100000... must have the same relative effect on the output as the switch from ...000000 to ...000001. The second problem is that even if the steady levels can be made correct, the transition process must result in no switching noise.

The idea that it is the *relative* effect that is important leads to reconsideration of delta modulation. Suppose, instead of trying to reconstruct an analog signal directly from a 16-bit data stream, we re-quantize the data first, in such a way as to reduce the number of bits required. Since 12-bit converters are relatively commonplace, we could consider reducing the 16-bit data to 12 bits, by truncating or approximating the 4 low-order bits. The first effect of our approximation will clearly be an increased quantization error, or reduced signal-to-noise ratio. The S/N ratio is decreased to 72 dB; better than a cassette player, but a long way short of the available quality. The problem is that we have produced only one approximated output number for each more precise input number. If—anticipating the result—we increase the sample rate by a factor of 2^4 , or 16, we would produce just over 700,000 samples per second from our CD data, instead of the input rate of about 44,000.


Figure C.1 illustrates the case of reducing the number of bits by 2. The approximation during the interval t_4 to t_8 of our $(N-2)$ -bit converter (Figure C.1(b)) to the input N -bit data (Figure C.1(a)) may be low when the signal is requantized with 2 bits fewer resolution. The next sample (t_8 to t_{12}) may be high, and the next low. In a high-frequency version, however, if the first sample is low the second can be forced high, as in Fig. C.1(c). The other outputs are chosen such that a 4-bit word is produced during the interval t_4 to t_8 to replace the single high-precision sample in such a way that the average level of the

output word is the same as that obtained by a high-precision conversion. The difference is that a relatively low-quality converter and a filter replace the costly 16-bit converter. (Note that we are not producing a digital word in the conventional sense to represent the original sample. In the 4-bit word of our high-rate output, the first data bit produced has the same weight as the last.) The average quantization error is zero, and the filtering process is simplified because much of the energy of the harmonics of the signal now occurs at 4 times the original sampling frequency.

(a) N-bit Input



(b) Output of (N-2)-bit Quantizer



(c) Output as in (b) with Higher Sampling Rate




Figure C.1 Requantization with higher sample rate

The process of deriving the requantized signal from the input is not particularly simple. Ultimately, at the cost of some memory use, and assuming fast enough memory hardware, it may be feasible to store the set of all possible output words in a memory that is addressed by the input data.

The method works because, in effect, the quantization noise is spread over a wider spectrum, so that the portion of it in the bandwidth of interest is greatly reduced. In the audio world, sampling frequencies of some tens of MHz are used with converters of a few bits (or perhaps only 1 bit) to produce outputs with dynamic range in the order of 100 dB.

REFERENCES

- Adolfsson, M., Einvall, C.H., Lindbergh, P., Samuelsson, J., Ahlgren, L. and Edlund, H. (1989). EHV series capacitor banks: A new approach to platform to ground signalling, relay protection and supervision. *IEEE Transactions on Power Delivery*, 4(2), 1369-1378.
- Agnew, P.G. (1911). A study of the current transformer with particular reference to iron loss. *Bulletin of the Bureau of Standards*, 7(3) 431.
- Ajjarapu, V. and Albanna, Z. (1991). Application of genetic based algorithms to optimal capacitor placement. *Proc. 1st International Forum on Applications of Neural Networks to Power Systems*. IEEE Pub. 91TH0374-9. NY: IEEE.
- American National Standards Institute. (1977). *IEEE standard dictionary of electrical and electronics terms*. ANSI/IEEE Std 100-1977.
- American National Standards Institute. (1978). *IEEE standard requirements for instrument transformers*. ANSI/IEEE C57.13-1978.
(The equivalent IEC documents are: Publication 185, Current Transformers; Publication 186, Voltage Transformers; Publication 44-3, Instrument Transformers, Part 3 Combined Transformers.)
- Annual Electric Utility Industry Forecast. (1991, October) *Electrical World*, pp. 11-22.
- Berkebile, L.E. (1980). Electronic current transducer for EHV circuits. (Report EL-1611, Final Report of Research Project 560-1). Performed by Westinghouse Electric Corporation: Electric Power Research Institute, Palo Alto CA.
- The Best Companies. (1991, September 7). *The Economist*, pp 21-24.
- Black, R.W., Formby, J.R. and Walker, R. (1985). The application of feeder automation to existing and future 11-kV urban networks. *CIREN: 8th International Conference on Electricity Distribution*, May 20-24, Brighton, UK.
- Cease, T.W. and Johnston, P.M. (1990). A magneto-optic current transducer. *IEEE Transactions on Power Delivery*. 5(2) 548-555.
- Control Center Survey: Installed and planned applications software. (1990, February). (Report EL-6634). Electric Power Research Institute, Palo Alto CA.
- Deliyannides, J.S., Millard, F.G., Kezunovic, M., Masterson, R.J., McGarry, M.J., Reuther, J.F., Sackin, M., Schwalenstocker, T.H., and Udren, E.A. (1981). Substation control and protection project: System requirements specification" (EL-1813 Interim Report of Research Project 1359-1). Performed by Westinghouse Electric Corporation: Electric Power Research Institute, Palo Alto CA.

- Distribution Automation. (1988). IEEE Working Group on Distribution Automation, (Eds.). (Tutorial Course text No. 88EH0280-8-PWR). New York: IEEE.
- Dolan, E.G. (1977). *Microeconomics*. Hinsdale, IL: The Dryden Press.
- Dougherty, J.J. and Kirkham, H. (1970). System aspects of tapped dc lines with dynamic controls. *IEEE Transactions on Power Apparatus and Systems*. 89(9), 2066-2074.
- ECAR's future: A bumpy road. (1991, June). *Electrical World*, 205(6), 21-22.
- Gould, S.J. (1982). *The Panda's Thumb*. New York: W.W. Norton.
- Hansen, D.H., Illian, H.F. and Kilar, L.A. (1991, May). Commonwealth Edison designs its own Scada/EMS system. *Electrical World*, 205(5), 35-38.
- Hebner, R.E. (1987) at an EPRI Workshop on the Role of Optical Sensors in Power System Voltage and Current Measurements, National Bureau of Standards, Gaithersburg, MD.
- Holman, J.P. (1978). *Experimental methods for engineers*. New York: McGraw-Hill.
- Holte, K.C. (1989, July). Technology requirements for a competitive electric utility market in the 21st century. *IEEE Power Engineering Review*, 9(7), 18-22.
- Houston, J.M. (1980). Electronic current transducer for high voltage direct current lines. (EL-1343 Final Report of Research Project 668-1). Performed by General Electric Company: Electric Power Research Institute, Palo Alto CA.
- Kendrew, T.J., and Marks, J.A. (1989). Automated distribution comes of age. *IEEE Computer Applications in Power*. 2(1), 7-10.
- Kirchmayer, L.K. (1958). *Economic Operation of Power Systems*. New York: John Wiley and Sons.
- Kirkham, H., Friend, H., Jackson, J. and Johnston, A. (1990). *Developments in fiber optics for distribution automation*. JPL Publication No. 90-2, Jet Propulsion Laboratory, Pasadena, California.
- Kirkham, H., Göttsche, A., Niebur, D., Friend, H., Johnston, A. (1991). *Communications and control for electric power systems: The 1990 report*. JPL Publication No. 91-44. Jet Propulsion Laboratory, Pasadena, California.
- Kirkham, H., Johnston, A.R. and Friend, H. (1989). *Distribution automation applications of fiber optics*. JPL Publication No. 89-10, Jet Propulsion Laboratory, Pasadena, California.
- Kron, G. (1951). Tensorial analysis of integrated transmission systems—Part I: The six basic reference frames. *AIEE Trans.* 70(1), 1197-1205.
- Lee, R.E. and Brooks, C.L. (1988). A method and its application to evaluate automated distribution control. *IEEE Transactions on Power Delivery*. 3(3), 1232-1240.
- MacDougall, T.W., Lutz, D.R. and Wandmacher, R.A. (1991) Development of a Fiber Optic Current Sensor for Power Systems. (PES Transmission and Distribution Conference and Exposition, Dallas, TX September 22-27 1991.) To be published in *IEEE Transactions*.
- Nara, K., Shiose, A., Kitagawa, M. and Ishihara, T. (1991) Implementation of genetic algorithm for distribution systems loss minimum re-configuration. Paper No. 91 SM 467-1 PWRs presented at the IEEE Power Engineering Society Summer Meeting in San Diego, CA, July 28-Aug 1. To be published in *IEEE Transactions*.
- National Energy Strategy. (1991). Stock No. 061-00-00754-7. Washington, DC: U.S. Government Printing Office.
- Pao, Y-H. (1990). Application of neural-net computing to power systems. *Proceedings of*

- NSF workshop on artificial neural-net methodology in power systems engineering, April 8–10. Washington, DC: National Science Foundation.**
- Papp, A. and Harms, H. (1980). Magneto-optical current transformer. *Applied Optics*, 19(22), 3729–3745.**
- Patton, J.B., Whittier, W.P. and Walk, S.R. (1991, November). Control of future energy management systems—Coping with uncertainty and complexity. *Proceedings of the workshop on real-time control and operation of electric power systems*. Available through Oak Ridge National Laboratory, Power Systems Technology Program, Oak Ridge, TN 37831.**
- Schweitzer, E.O., To, H.-W. and Ando, M. (1980). Digital revenue metering algorithm: Development, analysis, implementation, testing and evaluation. (EL-1601 Final Report of Research Project 1359-3). Performed by Ohio University: Electric Power Research Institute, Palo Alto CA.**
- So, E. (1987). A current-comparator-based wattmeter for high-voltage power measurements at very low power factors. Paper 87 WM 182-9 at the IEEE PES 1987 Winter Meeting.**
- So, E. (1990). The application of the current-comparator technique in instrumentation and measurement equipment for the calibration of non-conventional instrument transformers with non-standard rated outputs. Paper 90 IC 565-2 PWRD at the IEEE PES International Power Meeting in India October 28–November 1 1990.**
- Survey of Telecommunications. (1991, October 5). *The Economist*.**
- Ulmer, E.A. Jr. (1990). A high-accuracy optical current transducer for electric power systems. *IEEE Transactions on Power Delivery*, 5(2), 892–898.**
- 1990 System Disturbances. (1991). Princeton, NJ: North American Reliability Council.**

